z/OS

Network File System Guide and Reference

SC26-7417-05

z/OS

Network File System Guide and Reference

SC26-7417-05

Note:
Before using this information and the product it supports, be sure to read the general information under "Notices" on page 397.

Sixth Edition, January 2006

This edition applies to Version 1 Release 7 of z/OS® (5694-A01), Version 1 Release 7 of z/OS.e (5655-G52), and to all subsequent releases and modifications until otherwise indicated in new editions.

This edition replaces SC26-7417-04.

IBM® welcomes your comments. A form for readers' comments may be provided at the back of this document, or you may address your comments to the following address.

International Business Machines Corporation
Department 55JA, Mail Station P384
2455 South Road
Poughkeepsie, NY 12601-5400
United States of America

FAX (United States & Canada): 1+845+432-9405
FAX (Other Countries): Your International Access Code +1+845+432-9405

IBMLink™ (United States customers only): IBMUSM10(MHVRCFS)
Internet e-mail: mhvrcfs@us.ibm.com
World Wide Web: http://www.ibm.com/servers/eserver/zseries/zos/webqs.html
Network File System (NFS) Web site: http://www.ibm.com/servers/eserver/zseries/zos/nfs/

If you would like a reply, be sure to include your name, address, telephone number, or FAX number.

Make sure to include the following in your comment or note:
- Title and order number of this document
- Page number or topic related to your comment

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

© Copyright International Business Machines Corporation 1991, 2006. All rights reserved.
US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Figures ix

Tables xi

About this document xiii
Required product knowledge xiii
Where to find more information xiii
Access to documents xiii
Access to softcopy documents on CD-ROM and DVD collections and the Internet xiv
How to look up message explanations with LookAt xv
How to read syntax diagrams xvi
Related protocol specifications xviii

Summary of changes xxi

Part 1. Introduction 1

Chapter 1. Introduction 3
Overview 3
z/OS UNIX files 4
z/OS UNIX advantages. 4
NFS protocol compliance 5
Conventional MVS data sets 5
Mounting MVS data sets onto a client mount point 5
Crossing between file systems–NFS server . . . 6
Creating conventional MVS data sets 6
Serializing and sharing data sets. 6
Server control files 7
Attributes data set 7
Exports data set 7
Mount handle data set 8
Log data set 8
Checklist data set. 8
| Tested clients for the z/OS NFS server. 9
| NFS protocol attributes for the z/OS NFS server . 9
Supported servers for the z/OS NFS client 9
WebNFS support 9
| NFS version 3 and version 4 with TCP/IP protocols 10
| TCP/IP version 6 11
Data transfer and conversion 11
Native ASCII support - NFS V2 and V3 11
User-specified port range support 12
| Dynamic addressing 12

Part 2. NFS User's Guide 13

Chapter 2. Creating conventional MVS data sets 15
Overriding data set creation attributes 15
Preparing to create an MVS file. 15
Naming MVS files 16
Creating physical sequential files 16
Creating direct access files 17
Creating PDSs and PDSEs 17
Creating a PDS or PDSE - mkdir dsntype(pds), dsntype(library) 17
Removing a PDS or PDSE - rm, rmdir 18
Accessing PDS or PDSE members 18
Updating or extending a PDS or PDSE member 19
Timing out while writing a PDS or PDSE member 19
Wildcard copy to a PDS or PDSE 19
Limitations of a PDS 19
Concurrent writes to a PDSE 19
Creating VSAM files 19
Exploiting SAM striped files. 20
| Exploiting large format data sets 21

Chapter 3. Using conventional MVS data sets 23
Special MVS considerations 23
Selection of an MVS data storage format . . . 23
File size determination and time stamps. . . . 24
Ownership and permissions 24
NFS version 2 and version 3 statelessness . . . 24
NFS version 4 state. 25
File reading and writing 26
Case sensitivity–maplower, nomaplower. . . . 27
Selection of text or binary processing modes–text, binary 27
Number representation 28
Access to migrated files–retrieve, noretrieve; wait, nowait 28
Access to migrated system-managed data sets . . 29
File handle refresh 29
Mapping between the workstation and MVS file systems. 30
File name extension mapping 30
Mounting of MVS data sets onto a client mount point 31
Use of a PDS or PDSE as a directory 33
Use of multiple mount points 33
Data set serialization and sharing 34
NFS protocol 34
NFS file system attributes 34

Chapter 4. Using z/OS UNIX System Services files 35
z/OS UNIX file system 35
POSIX compatibility 36
NFS protocol 36
Attributes specific to z/OS UNIX System Services 36
Synchronous write to a z/OS UNIX file for NFS version 2 protocol 37

© Copyright IBM Corp. 1991, 2006

Synchronous write to a z/OS UNIX file for NFS version 3 or 4 protocol. 37
Authorization checking when writing to a z/OS UNIX file 37
HFS site attribute 37
Protecting your z/OS UNIX System Services files . 38
Accessing z/OS UNIX System Services files from a client 38
Mount examples. 38
z/OS UNIX data transfer and conversion 39
Data transfer under the NFS version 4 protocol 39
Text or binary processing - NFS version 2 and 3 protocols 39
Linking an MVS data set to a z/OS UNIX file system 40
Creating an external link 40
Displaying the contents on an external link. . . 41
Deleting an external link 41
UNIX look and feel. 41
NFS file system attributes 42

Chapter 5. Locking and access control 43
Using Network Lock Manager (NLM) in NFS V2 and V3 43
Monitored locks 43
Non-monitored locks 43
Specifying a grace period for reclaiming locks . . 44
Listing locks held for a file 44
Releasing locks held for a file 44
Using Network Status Monitor (NSM) in NFS V2 and V3 44

Chapter 6. Commands and examples for AIX and UNIX clients 47
Using commands on AIX 47
Quick reference of AIX and UNIX commands . . . 49
Accessing z/OS UNIX System Services and conventional MVS files 50
Mvslogin command examples 51
Mount command examples 52
Using mount to override server default attributes 53
Displaying default and mount point attributes - showattr 54
Unmounting and logging out of z/OS 57
Disconnecting your mount point - umount . . . 57
Ending your z/OS session - mvslogout 58

Chapter 7. Commands and examples for z/OS NFS clients 59
Using commands on DFSMS/MVS 59
Accessing z/OS 63
Mvslogin command examples 64
Mount command syntax and examples 65
Unmount command syntax and examples . . . 71
Displaying client statistical information–nfsstat 72
Displaying server mount information–showmount 75
Displaying default and mount point attributes–showattr 76
Ending your z/OS session - mvslogout 77

Chapter 8. Initialization attributes for the z/OS NFS client 79
Client attribute syntax 79
Datacaching attribute 84
Mount processing parameters and installation parameters 84
NFS client translation support 85
z/OS NFS client with z/OS NFS server 86

Chapter 9. Initialization attributes for the z/OS NFS server 87
Attributes used for z/OS UNIX System Services file access 87
Multipliers 87
Duplicate attributes. 88
Data set creation attributes syntax 88
Processing attributes syntax 93
Timeout attributes 99
Retrieve attributes 99
Mapped keyword processing attribute 100
Native ASCII processing attributes 100
Site attributes syntax 104

Part 3. Customization and Operations 117

Chapter 10. Customization 119
Protecting your programs and files 119
Protecting the server control files 119
Setting up the z/OS NFS authorization 119
Protecting the file system on z/OS with the NFS V4 protocol 120
Protecting the file system on z/OS with the Security site attribute 123
Customizing installation security exits 126
Using UNIX style credentials for authentication 126
Converting data 126
Creating the conversion environment for Unicode Services 127
Configuring NFS servers for multiple TCP/IP stacks 129
Invocation 129
Example procedure to start an NFS server in a multiple server environment 129
User interactions 130
Errors 130
Messages and codes 130
Collecting NFS usage data 130
Configuring the z/OS NFS client 131
Creating the PARMLIB statement for the client 132
Updating z/OS system data sets for the client 132
Allocating client log data sets 133
Mounting remote file systems 133
Configuring the z/OS NFS server 133
Attributes data set. 133
Exports data set 134
Checklist data set 138
Mount handle data sets 139
Lock data sets 139

Converting data between ASCII and EBCDIC - NFS V2 and V3 only 140
Updating z/OS system data sets for the server 142
Allocating the z/OS NFS server log data sets 142
Allocating and modifying mapping side file . . 142
Modifying tcpip.ETC.RPC 142
Setting up a user specified port range 143
Configuring a secure z/OS NFS server 144
Using dynamic client IP addressing 147
Installing the client enabling commands 147
Retrieving commands for AIX, Sun Solaris, and Linux 150
Porting the mvslogin, mvslogout, and showattr commands 153
Porting on different compilers and operating systems 154

Chapter 11. Network File System operation 157
Starting the z/OS NFS client 157
Stopping the z/OS NFS client 158
Starting the z/OS NFS server 158
Starting multiple servers 160
Stopping the z/OS NFS server 160
Starting the z/OS NFS NSM and z/OS NFS NLM 161
Starting component tracing for the z/OS NFS server 161
Entering operands of the modify command for the z/OS NFS server 164
Exportfs operand 166
Freeze operand 166
List operand 166
Mapfile operand 167
Release operand 167
Status operand 168
Unmount operand 169
Unmntall operand 169
Unmnthfs operand 169
Unmntmvs operand 169
Entering operands of the modify command for diagnosis 170
Flushctr operand 170
Flushlog operand 170
Listlock operand 170
Log=msglevel operand 171
Smf operand 172
Version operand 172

Chapter 12. Installation-wide exits for the z/OS NFS server 173
Requirements for NFS 173
Sample link-edit JCL 174
Storage blocks of the server exits 174
Login installation-wide exit 175
Requirements of the login exit 177
Options of the login exit 177
Structure of the login exit message 177
Contents of the login exit parameter list . . . 177
Login exit parameter list 178
Request codes to the login exit 178
Return codes from the login exit 179
System initialization routine of the login exit 179
Start of new user session routine of the login exit 179
User login request routine of the login exit . . 180
User logout request routine of the login exit . . 180
System termination routine of the login exit . . 181
File security installation-wide exit 181
Requirements of the file security exit 184
Structure of the file security exit message . . . 184
Contents of the file security exit parameter list 184
File security exit parameter list 186
Request codes to the file security exit 186
Return codes from the file security exit 186
Validate allocate request routine of the file security exit 186
Validate write request routine of the file security exit 187
Validate read request routine of the file security exit 187
Return security permissions routine of the file security exit 188

Part 4. Performance Tuning 191

Chapter 13. Performance tuning in the NFS environment 193
What is performance tuning? 193
How is performance characterized? 193
What is the NFS environment? 194
How to tune for performance 194
Impact of the NFS protocol on performance . . . 196

Chapter 14. Optimizing the NFS environment 199
Network performance tuning 199
NFS client system performance tuning 200
NFS server system performance tuning 201
z/OS constraints 202

Chapter 15. Evaluating z/OS NFS performance 205
Evaluating throughput 205
Single process throughput 205
Multiple process throughput 206
Multiple client throughput 206
Evaluating NFS command response time 206
Evaluating CPU utilization 206
Collecting server usage data 207

Chapter 16. Tuning the z/OS NFS server 209
Data set creation attributes 209
Block size and record length 209
Record format 210
Data set organization and data set type. . . . 210
Processing attributes 211
Character translation 211
File size determination 211

Data set timeout specification 213
Accessing migrated files 213
Asynchronous z/OS UNIX processing 214
Site attributes 214
Buffer usage and caching 214
Ordering out-of-sequence data. 215
Storage considerations 217
Subtasking 217

Chapter 17. Tuning the z/OS NFS client 219
Caching 219
Dynamicsizeadj. 219
Bufhigh 220
Biod 220
Readahead 220
Delaywrite 220
Vers 220
Wsize and rsize 221

Part 5. Diagnosis and Messages 223

Chapter 18. Diagnosis and reporting of problems 225
Correcting input errors 225
Using keywords to identify a problem 226
Component identification keyword 226
Release level keyword 226
Type-of-failure keyword 227
Service level keyword 229
Using z/OS component tracing 229
Component trace benefits 230
Using NFS component trace PARMLIB member CTINFS*nn* 230
Capturing NFS component trace information in an SVC dump 232
Using a z/OS component trace external writer 233
Setting up a dump data set for abnormal ends . . 235
Searching the IBM database for APARs and PTFs 235
Contacting the IBM Support Center 235
Diagnostic aids 236
First failure data capture 236
Errors and messages 237
Debug trace data capture 238
Environmental checklist 238

Chapter 19. Network File System messages 241
Server messages 241
Client messages 281
Messages from the client platform (AIX) 295

Chapter 20. Return codes 299

Chapter 21. Reason codes 309

Appendix A. File size value for MVS data sets 313
Storage of the file size value 313
System-managed PS, VSAM, and PDSE data sets. 313
Migrated system-managed data sets 313
Non-system-managed, PDS, and direct access data sets 314
How the file size value is generated 314
Using fastfilesize to avoid read-for-size 314
Using nofastfilesize 315

Appendix B. Time stamps for MVS data sets 317
Time stamps for system-managed VSAM and PS data sets 317
Time stamps for non-system-managed PS and DA data sets 317
Storage of time stamps 317
Client program requirements 318
Generating time stamps 318
Time stamps for non-system-managed VSAM data sets. 318
Time stamps for PDSs and PDSEs 318
Setting time stamps 320

Appendix C. NFS attributes 321
NFS file system attributes for MVS data sets . . . 321
NFS file system attributes for z/OS UNIX file systems 322
NFS protocol attributes 323

Appendix D. NSM (statd) protocol . . 327
Using supported NSM (statd) procedures 327

Appendix E. Sample CHECKLIST data set 329

Appendix F. NFS system server sample attribute table. 331

Appendix G. Sample exports data set 349

Appendix H. Sample startup procedures 353
Sample z/OS NFS server startup procedures . . . 353
Sample z/OS NFS client startup procedures . . . 356

Appendix I. Retrieval of source code for client enabling commands 359

Appendix J. PCNFSD protocol 361
Accessing data with PCNFSD 361
Accessing z/OS UNIX files. 361
Starting the PCNFSD server 361
Using supported PCNFSD protocols. 362
Version 1 of the PCNFSD protocol 362
Version 2 of the PCNFSD protocol 362

Appendix K. SMF C and assembler header macros 365
SMF C header macro GFSASSMF. 365
SMF assembler header macro GFSAUSMF. 372

Appendix L. Capturing diagnostic information using z/OS NFS log data sets and from other components . . . 379
Using log data sets 379
Server log data sets 379
Client log data sets 380
Debug trace data capture 381
z/OS NFS server debug trace capture 381
z/OS NFS server DEBUG trace types 382
z/OS NFS client debug trace capture 383
Related component trace capture 384
z/OS UNIX System Services activity trace . . . 384
z/OS Hierarchical File System (HFS) Physical File System activity trace 384
z/OS TCP/IP activity trace. 384
AIX client activity trace 385
SUN client activity trace. 386
z/OS dump 386

Appendix M. GFSAMHDJ sample code for creating NFS mount handle data sets and lock data sets 387

Appendix N. Setting up a Kerberos Key Distribution Center 391

Appendix O. Accessibility 395
Using assistive technologies 395
Keyboard navigation of the user interface 395
z/OS information 395

Notices 397
Programming interface information 398
Trademarks 398

Glossary 399

Index 407

Figures

1. NFS client-server relationship 3
2. Examples of mounting MVS data sets on Windows, UNIX and Linux clients 33
3. Example of mounting an HFS or zFS file from a UNIX client 35
4. Syntax for commands 47
5. Displaying default attributes 55
6. Displaying default and mount point attributes 56
7. TSO MOUNT command syntax operands 65
8. TSO UNMOUNT command syntax operands 71
9. Displaying NFS client rpc and NFS statistical information 73
10. Displaying NFS client rpc statistical information 73
11. Displaying NFS client NFS statistical information 74
12. Displaying NFS mounted file system information 74
13. Displaying V.2 protocol NFS mounted file system information 74
14. Displaying version 3 protocol NFS mounted file system information. 75
15. Displaying default attributes 77
16. Sample filesystype parmlib statement 132
17. Entry for a directory 136
18. Specifying the mount handle data set in the MVSNFS procedure 139
19. Specifying the lock data set in the MVSNFS procedure 140
20. Modify /etc/services for mountd, mvsmount, pcnfsd, and showattr 143
21. Modify tcpip.profile for NLM integration 143
22. Retrieving the client enabling commands for AIX, Sun Solaris, and Linux. 150
23. Common source files 154
24. Summary of the modify command 164
25. Sample link-edit JCL for the NFS 174
26. Determining which login checking routines are used 176
27. Determining which file security checking routines are used 183
28. Displaying NFS client rpc and NFS statistical information 197
29. Sample network topology 199
30. Directory list comparison between DFSMS-managed and non-managed 212
31. Bufhigh utilization with percentsteal 214
32. Relationship between cachewindow and BIODs 216
33. Logicalcache utilization for cachewindows 216
34. z/OS NFS Server component trace PARMLIB member CTINFS00. 232
35. NFS system server sample attribute table 331
36. Sample exports data set 349
37. Sample z/OS NFS server startup procedures 353
38. Sample z/OS NFS client startup procedures 356
39. Retrieving source code for client enabling commands 359
40. SMF C header macro GFSASSMF 365
41. SMF assembler header macro GFSAUSMF 372
42. Sample code for creating mount handle data sets and lock data sets 387

© Copyright IBM Corp. 1991, 2006

Tables

1. Reference documents xiii
2. View of NFS server capability 10
3. Breakdown of text and binary writes 27
4. Examples of mounting MVS data sets from clients 31
5. NFS procedures 34
6. Examples of the mount command for clients 38
7. Examples of the mvslogin command for clients 51
8. Examples of the showattr command for clients 54
9. Examples of the umount command for clients 57
10. Example of the mvslogout command for clients 58
11. Examples of the mvslogin command for clients 64
12. Attributes - z/OS NFS client 79
13. Client attributes 79
14. Mount processing parameters 84
15. Installation parameters 85
16. z/OS NFS clients with non-z/OS based NFS servers 85
17. Attributes - z/OS NFS server 87
18. Data set creation attributes 88
19. Processing attributes 93
20. The mapped keyword and existing keywords 100
21. File tagging with Unicode Services active 101
22. File tagging with Unicode Services not active 103
23. Site attributes 104
24. Customizing NFS 119
25. z/OS server processing of a mount request 125
26. z/OS server processing of a file request 125
27. Modifying tcpip.ETC.RPC 142
28. Modify /etc/rpc for mvsmount and showattr 143
29. Files in the prefix.NFSTARB data set to download to clients 149
30. List of installation-wide exits 173
31. Format of login installation-wide exit routine parameter list 178
32. Request codes to the login exit 178
33. Return codes from the login exit 179
34. Codes and fields for system initialization 179
35. Codes and fields for start of new user session 179
36. Codes and fields for user login request 180
37. Codes and fields for logout request 180
38. Codes and fields for system termination 181
39. Format of the parameter list for the file security installation-wide exit 185
40. Request codes to the file security exit 186
41. Return codes from the file security exit 186
42. Codes and fields for validate allocate request 186
43. Codes and fields for validate write request 187
44. Codes and fields for validate read request 187
45. Codes and fields for return security permissions 188
46. Network performance tuning symptom and action information 199
47. NFS client system performance tuning symptom and action information 200
48. z/OS constraints symptom and action information 202
49. Default block sizes for RAMAC 3 or ESS 2105 DASD 210
50. Client installation parameters for tuning 219
51. Mount parameters for performance and tuning 219
52. Summary of type-of-failure keywords 227
53. NFS symptom data 236
54. Dump content and storage areas 237
55. Diagnostic errors and messages 238
56. Messages - client operating system, NFS server, and NFS client. 241
57. Message format for the NFS server log data set 241
58. NFS client z/OS operators console message format 281
59. NFS client log data set message format 281
60. Common variables 282
61. Externalized return codes defined by the NFS version 2 protocol 299
62. Externalized return codes defined by the NFS version 3 protocol 299
63. Externalized return codes defined by the NFS version 4 protocol 300
64. z/OS NFS Server: z/OS UNIX return codes mapped to NFS Version 2 return codes . . . 302
65. z/OS NFS Server: z/OS UNIX return codes mapped to NFS Version 3 return codes . . . 304
66. z/OS NFS Server: z/OS UNIX return codes mapped to NFS Version 4 return codes . . . 306
67. Parsing error (when reason code is 6E01xxxx) 309
68. Parsing error (when reason code is from 6E0111yy to 6E0129yy) 310
69. Parsing error (when reason code is 6Exxyyyy') 311
70. Time stamp sources for PDS and PDSE members 319
71. Time stamp sources for PDS and PDSE data sets (directories) 319
72. File system values to get dynamic file system information 321
73. File system values to get static file system information 321
74. File system values to retrieve POSIX information 321
75. File system values to get static file system information 322
76. File system values to retrieve POSIX information 322
77. NFS Version 4 Attributes 323

© Copyright IBM Corp. 1991, 2006

About this document

This document provides users, system programmers, and operators with information about using, customizing, operating, tuning, and diagnosing the z/OS® Network File System (z/OS NFS).

For information about the accessibility features of z/OS, for users who have a physical disability, see Appendix O, "Accessibility," on page 395.

Required product knowledge

To use this document effectively, you should be familiar with the IBM multiple virtual system (MVS) as a component of the z/OS operating system, the IBM Time Sharing Option (TSO), and their commands. In addition, you should be familiar with System Modification Program/Extended (SMP/E) and the basic concepts of the NFS protocol and networking (Transmission Control Protocol/Internet Protocol (TCP/IP)).

Where to find more information

Where necessary, this document references information in other documents, using the shortened version of the document title. For complete titles and order numbers of the documents for all products that are part of z/OS, see *z/OS Information Roadmap*.

Access to documents

Table 1 contains additional reference information.

Table 1. Reference documents

Title	Order Number
AIX Commands Reference, Volume 1	SC23-2537
AIX Commands Reference, Volume 2	SC23-2538
AIX Commands Reference, Volume 3	SC23-2539
AIX Commands Reference, Volume 4	SC23-2539
AIX Commands Reference, Volume 5	SC23-2639
AIX General Concepts and Procedures for RISC System/6000	GC23-2202
Character Data Representation Architecture Overview	GC09-2207
Character Data Representation Architecture Reference and Registry	SC09-2190
SMP/E Reference	SA22-7772
SMP/E User's Guide	SA22-7773
z/OS and Software Products DVD Collection	SK3T-4271
z/OS Collection	SK3T-4269
z/OS Communications Server: IP Configuration Reference	SC31-8776
z/OS DFSMS Introduction	SC26-7397
z/OS DFSMS Macro Instructions for Data Sets	SC26-7408
z/OS Migration	GA22-7499
z/OS DFSMS Using the New Functions	SC26-7473
z/OS DFSMS Using Data Sets	SC26-7410

© Copyright IBM Corp. 1991, 2006

Table 1. Reference documents (continued)

Title	Order Number
z/OS DFSMSdfp Advanced Services	SC26-7400
z/OS DFSMSdfp Diagnosis	GY27-7618
z/OS DFSMSdfp Storage Administration Reference	SC26-7402
z/OS DFSMShsm Diagnosis	GC52-1083
z/OS Information Roadmap	SA22-7500
z/OS MVS Installation Exits	SA22-7593
z/OS MVS JCL Reference	SA22-7597
z/OS MVS Programming: Authorized Assembler Services Guide	SA22-7608
z/OS MVS Programming: Authorized Assembler Services Reference ALE-DYN	SA22-7609
z/OS MVS Programming: Authorized Assembler Services Reference ENF-IXG	SA22-7610
z/OS MVS Programming: Authorized Assembler Services Reference LLA-SDU	SA22-7611
z/OS MVS Programming: Authorized Assembler Services Reference SET-WTO	SA22-7612
z/OS MVS System Codes	SA22-7626
z/OS MVS Programming: Authorized Assembler Services Guide	SA22-7608
z/OS MVS System Management Facilities (SMF)	SA22-7630
z/OS MVS System Messages, Vol 1 (ABA-AOM)	SA22-7631
z/OS MVS System Messages, Vol 2 (ARC-ASA)	SA22-7632
z/OS MVS System Messages, Vol 3 (ASB-BPX)	SA22-7633
z/OS MVS System Messages, Vol 4 (CBD-DMO)	SA22-7634
z/OS MVS System Messages, Vol 5 (EDG-GFS)	SA22-7635
z/OS MVS System Messages, Vol 6 (GOS-IEA)	SA22-7636
z/OS MVS System Messages, Vol 7 (IEB-IEE)	SA22-7637
z/OS MVS System Messages, Vol 8 (IEF-IGD)	SA22-7638
z/OS MVS System Messages, Vol 9 (IGF-IWM)	SA22-7639
z/OS MVS System Messages, Vol 10 (IXC-IZP)	SA22-7640
z/OS Security Server RACF Security Administrator's Guide	SA22-7683
z/OS Security Server RACF System Programmer's Guide	SA22-7681
z/OS Security Server RACF Callable Services	SA22-7691
z/OS Integrated Security Services Network Authentication Service Administration	SC24-5926
z/OS Integrated Security Services Network Authentication Service Programming	SC24-5927
z/OS TSO/E User's Guide	SA22-7794
z/OS Support for Unicode: Unicode Services	SA22-7649
z/OS UNIX System Services File System Interface Reference	SA22-7808
z/OS UNIX System Services Messages and Codes	SA22-7807
z/OS UNIX System Services Programming: Assembler Callable Services Reference	SA22-7803
z/OS UNIX System Services Programming Tools	SA22-7805
z/OS UNIX System Services Command Reference	SA22-7802
z/OS UNIX System Services User's Guide	SA22-7801
z/OS UNIX System Services Planning	GA22-7800

Access to softcopy documents on CD-ROM and DVD collections and the Internet

This book will also be available on the following collections the next time they are updated.

z/OS V1Rx Collection, SK3T-4269

z/OS V1Rx and Software Products DVD Collection, SK3T-4271

This book will also be available in BookManager® and PDF format on the z/OS Internet library at:

`http://www.ibm.com/servers/eserver/zseries/zos/bkserv/`

Using a BookManager READ program, you can view and find information quickly in BookManager documents in a variety of environments. You can view the books directly from a CD-ROM or DVD, or copy the books to a shared workstation or local area network (LAN) server. You can also transfer the books to your host system or view them on the Internet. For instance, the z/OS product includes both BookManager READ for host viewing and BookManager BookServer, which allows you to access and read books over an Internet or intranet connnection using an HTML browser. From CD-ROM or DVD you can use any supported IBM BookManager reader, such as the IBM Softcopy Reader. For more information, see:

`http://www.ibm.com/servers/eserver/zseries/softcopy/`

You can view and print PDF files using an Adobe Acrobat Reader available free on the Internet at:

`http://www.adobe.com/prodindex/acrobat/`

How to look up message explanations with LookAt

LookAt is an online facility that lets you look up explanations for most of the IBM® messages you encounter, as well as for some system abends and codes. Using LookAt to find information is faster than a conventional search because in most cases LookAt goes directly to the message explanation.

You can use LookAt from the following locations to find IBM message explanations for z/OS elements and features, z/VM®, and VSE:

- The Internet. You can access IBM message explanations directly from the LookAt Web site at http://www.ibm.com/eserver/zseries/zos/bkserv/lookat/.
- Your z/OS TSO/E host system. You can install code on your z/OS or z/OS.e systems to access IBM message explanations, using LookAt from a TSO/E command line (for example, TSO/E prompt, ISPF, or z/OS UNIX® System Services running OMVS).
- Your Windows® workstation. You can install code to access IBM message explanations on the *z/OS Collection* (SK3T-4269), using LookAt from a Windows DOS command line.
- Your wireless handheld device. You can use the LookAt Mobile Edition with a handheld device that has wireless access and an Internet browser (for example, Internet Explorer for Pocket PCs, Blazer, or Eudora for Palm OS, or Opera for Linux® handheld devices). Link to the LookAt Mobile Edition from the LookAt Web site.

You can obtain code to install LookAt on your host system or Windows workstation from a disk on your *z/OS Collection* (SK3T-4269), or from the LookAt Web site (click **Download**, and select the platform, release, collection, and location that suit your needs). More information is available in the LOOKAT.ME files available during the download process.

How to read syntax diagrams

Throughout this library, diagrams are used to illustrate the programming syntax. Keyword parameters are parameters that follow the positional parameters. Unless otherwise stated, keyword parameters can be coded in any order. The following list tells you how to interpret the syntax diagrams:

- Read the diagrams from left-to-right, top-to-bottom, following the main path line. Each diagram begins on the left with double arrowheads and ends on the right with two arrowheads facing each other.

- If a diagram is longer than one line, each line to be continued ends with a single arrowhead and the next line begins with a single arrowhead.

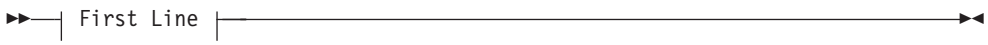

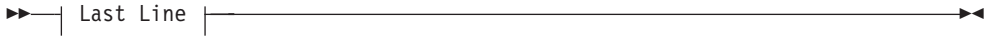

- Required keywords and values appear on the main path line. You must code required keywords and values.

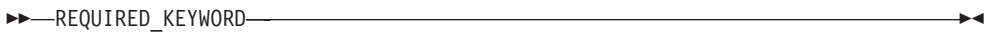

  If several mutually exclusive required keywords or values exist, they are stacked vertically in alphanumeric order.

- Optional keywords and values appear below the main path line. You can choose not to code optional keywords and values.

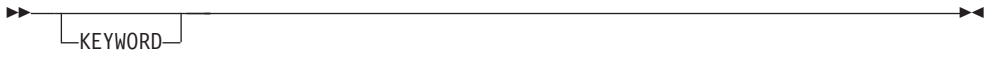

  If several mutually exclusive optional keywords or values exist, they are stacked vertically in alphanumeric order below the main path line.

- An arrow returning to the left above a keyword or value on the main path line means that the keyword or value can be repeated. The comma means that each keyword or value must be separated from the next by a comma.

- An arrow returning to the left above a group of keywords or values means more than one can be selected, or a single one can be repeated.

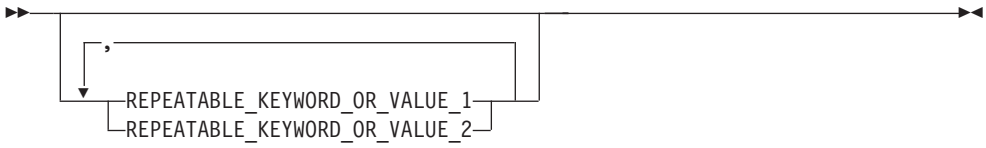

- A word in all uppercase is a keyword or value you must spell exactly as shown. In this example, you must code ***KEYWORD***.

  If a keyword or value can be abbreviated, the abbreviation is discussed in the text associated with the syntax diagram.
- If a diagram shows a character that is not alphanumeric (such as parentheses, periods, commas, and equal signs), you must code the character as part of the syntax. In this example, you must code ***KEYWORD=(001,0.001)***.

- If a diagram shows a blank space, you must code the blank space as part of the syntax. In this example, you must code ***KEYWORD=(001 FIXED)***.

- Default keywords and values appear above the main path line. If you omit the keyword or value entirely, the default is used.

- A word in all lowercase italics is a *variable*. Where you see a variable in the syntax, you must replace it with one of its allowable names or values, as defined in the text.

- References to syntax notes appear as numbers enclosed in parentheses above the line. Do not code the parentheses or the number.

Notes:

1 An example of a syntax note.

- Some diagrams contain *syntax fragments*, which serve to break up diagrams that are too long, too complex, or too repetitious. Syntax fragment names are in mixed case and are shown in the diagram and in the heading of the fragment. The fragment is placed below the main diagram.

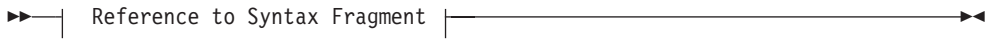

Syntax Fragment:

Related protocol specifications

IBM is committed to industry standards. The internet protocol suite is still evolving through Requests for Comments (RFC). New protocols are being designed and implemented by researchers, and are brought to the attention of the internet community in the form of RFCs. Some of these are so useful that they become a recommended protocol. That is, all future implementations for TCP/IP are recommended to implement this particular function or protocol. These become the *de facto* standards on which the TCP/IP protocol suite is built.

The Network File System (NFS) is implemented as a set of RPC procedures that use External Data Representation (XDR) encoding to pass arguments between clients and servers. The NFS is based on the following RFCs.

Internet Protocol	RFC 791, J.B. Postel
NFS: Network File System Version 2 Protocol Specification	RFC 1094, Sun Microsystems, Incorporated
NFS: Network File System Version 3 Protocol Specification	RFC 1813, Sun Microsystems, Incorporated
NFS: Network File System Version 4 Protocol Specification	RFC 3530, Sun Microsystems, Incorporated
Open Group Technical Standard Protocols for Interworking: XNFS, Version 3W	Document Number: C702
RPC: Remote Procedure Call Protocol Specification Version 2	RFC 1057, Sun Microsystems Incorporated
RPC: Remote Procedure Call Protocol Specification Version 2	RFC 1831, R. Srinivasan
User Datagram Protocol	RFC 768, J.B. Postel
WebNFS Client Specification	RFC 2054, B. Callaghan
WebNFS Server Specification	RFC 2055, B. Callaghan
XDR: External Data Representation Standard	RFC 1014, Sun Microsystems, Incorporated
XDR: External Data Representation Standard	RFC 1832, R. Srinivasan
Generic Security Service Application Program Interface, Version 2	RFC 2078, J.Linn, OpenVision Technologies
RPCSEC_GSS Protocol Specification	RFC 2203, M. Eisler, A.Chiu, L. Ling
The Kerberos Version 5 GSS-API Mechanism	RFC 1964, J.Linn, OpenVision Technologies
The Kerberos Network Authentication Service (V5)	RFC 1510, J. Kohl, Digital Equipment Corporation, C. Neuman, ISI

For more information about Request for Comments (RFC), see the Internet Engineering Task Force (IETF) home page:

`http://www.ietf.org/`

Summary of changes

This document contains terminology, maintenance, and editorial changes. Technical changes or additions to the text and illustrations are indicated by a vertical line to the left of the change.

You may notice changes in the style and structure of some content. For example, text that has a different look and format. The changes are ongoing improvements to the consistency and retrievability of information in our documents.

Summary of changes for SC26-7417-05, z/OS Version 1 Release 7

This document contains information that was previously presented in *z/OS Network File System Guide and Reference*, SC26-7417-04. The following sections summarize the changes to that information.

New Information

This edition includes new information about NFS version 4 security enhancements that were made available with PTF UA23455. See especially the following new sections:

- "Configuring a secure z/OS NFS server" on page 144
- Appendix N, "Setting up a Kerberos Key Distribution Center," on page 391.

Changed information

This edition includes changed information:

- Revised mount examples in Table 4 on page 31 and Table 6 on page 38.
- Expanded text and examples in "Protecting the file system on z/OS with the NFS V4 protocol" on page 120 and "Exports data set" on page 134.

Deleted information

This edition deleted the following information:

- Obsolete references to the 3172 controller and TCP/IP offload have been deleted from Chapter 14, "Optimizing the NFS environment," on page 199.
- Messages GFSA341I, GFSA342I, and GFSA343I have been deleted.

Summary of changes for SC26-7417-04, z/OS Version 1 Release 7

This document contains information that was previously presented in *z/OS Network File System Customization and Operation*, SC26-7417-03; *z/OS Network File System User's Guide*, SC26-7419-03; and *z/OS Network File System Performance Tuning Guide*, SC26-7418-03. The following sections summarize the changes to that information.

Some actions need to be taken before using the Network File System (NFS) in the z/OS V1R7 release; for information on those actions, see the NFS section of *z/OS Migration*.

New information

This edition includes new information:

- The z/OS NFS server's support for the NFS version 4 protocol has been added.

© Copyright IBM Corp. 1991, 2006

- NFS version 4 security enhancements to the z/OS NFS server are being made available with APAR OA11875. For details on these enhancements, see "Protecting the file system on z/OS with the NFS V4 protocol" on page 120.
- Component trace support for the NFS server has been added to "Using z/OS component tracing" on page 229 and Chapter 11, "Network File System operation," on page 157.
- A lease interval and other locking functions from the NFS version 4 protocol are added to Chapter 5, "Locking and access control," on page 43. The new site attribute **leasetime(n)**, which sets a time limit for clients to retrieve locks after a server restart, has been added to Table 23 on page 104.
- The following site attributes have also been added to Table 23 on page 104:
  - **nlm | nonlm**, for starting NLM and NSM as part of NFS initialization
  - **dhcp**, which allows dynamic client IP addresses
  - **denyrw**, which allows NFS V4 deny mode requests
  - **mvssec**, **hfssec**, and **pubsec**, which enforce client security access to file systems through the z/OS NFS server.
  - **fileidsize(32) | fileidsize(64)**, which defines the size of the file ID attribute that the z/OS NFS server returns.
- The following client attributes have been added to Table 13 on page 79:
  - **secure(udp)**, to specify the transport protocol for the NFS client
  - **disablella**, to disable Lookup Look-Aside caching.
- The processing attribute **mvsmnt** has been added to "Processing attributes syntax" on page 93.
- The NFS client mount processing parameters **rpcbind** and **convserv** have been added to "Mount processing parameters and installation parameters" on page 84.
- Information on accessing large format data sets has been added to "Exploiting large format data sets" on page 21.
- New fields for Internet Protocol version 6 addresses have been added to "Login installation-wide exit" on page 175 and "File security installation-wide exit" on page 181, and SMF records as described in "Collecting NFS usage data" on page 130.
- Lock data sets must be reallocated as VSAM data sets, as shown in "Lock data sets" on page 139.
- Mount handle data sets must be reallocated with larger record lengths to support IP version 6 and dynamic IP addresses, as shown in "Mount handle data sets" on page 139.
- Version commands, for determining the precise release and maintenance levels of NFS server and client code, have been added to Chapter 18, "Diagnosis and reporting of problems," on page 225, "Entering operands of the modify command for the z/OS NFS server" on page 164 and Chapter 7, "Commands and examples for z/OS NFS clients," on page 59.
- New recommended port range reservations are added in "Setting up a user specified port range" on page 143.
- New messages have been added to Chapter 19, "Network File System messages," on page 241.

Changed information

This edition includes changed information:

- References to the MVS™ operating system have been changed to "z/OS."

- References to “OpenEdition” and many references to “HFS” have been changed to ″z/OS UNIX.”

Moved information

This edition includes moved information:

- Information from the three previous z/OS NFS books has been organized into separate Parts of this publication.
- Reason code information has been moved into a separate chapter, Chapter 21, “Reason codes,” on page 309.

Deleted information

This edition deleted the following information:

- Procedures for manually starting and stopping the NFS Network Status Monitor (NSM) and Network Lock Manager (NLM) have been removed; they are replaced by the **nlm | nonlm** site attribute.
- The **nohfs** site attribute value has been deleted. It was unsupported starting in z/OS V1R5 and if specified, is ignored.
- Messages GFSA323I, GFSA324I, GFSA336I, GFSA337I, GFSA338I, GFSA344I, GFSA345I, GFSA565I, and GFSA783 have been deleted.

Summary of changes for SC26-7417-03, z/OS Version 1 Release 5

This document contains information that was previously presented in *z/OS Network File System Guide and Reference*, SC26-7417-02. The following sections summarize the changes to that information.

New information

This edition includes new information:

- Support for multivolume data sets.
- Support; see “User-specified port range support” on page 12 and “Setting up a user specified port range” on page 143.
- Added “Processing attributes syntax” on page 93; see **tag** attribute.
- Added parameter information and additional “Site attributes syntax” on page 104.
- Added “Mapfile operand” on page 167.
- Additional operands; “Unmntall operand” on page 169, “Unmnthfs operand” on page 169, and “Unmntmvs operand” on page 169.
- Error handling support;
  - “Site attributes syntax” on page 104.
  - Chapter 19, “Network File System messages,” on page 241.
- Messages; see “Server messages” on page 241.
- Error information; see Table 69 on page 311.
- Server support for the additional procedure from PCNFSD protocol; see “Procedure 12: perform mapping” on page 362.
- Added user information; see Appendix O, “Accessibility,” on page 395.

Changed information

This edition includes changed information:

- Attribute options, see “Setting up the z/OS NFS authorization” on page 119.

- Commands, see “Installing the client enabling commands” on page 147 and “Porting the mvslogin, mvslogout, and showattr commands” on page 153.
- Server options; see Table 21 on page 101.
- Procedure name messages; see Chapter 19, “Network File System messages,” on page 241.
- Sample data set; see Appendix E, “Sample CHECKLIST data set,” on page 329.
- Server sample; see Appendix F, “NFS system server sample attribute table,” on page 331.
- Server sample; see Figure 37 on page 353.
- Client sample; see “Sample z/OS NFS client startup procedures” on page 356.
- Data set file; see Appendix G, “Sample exports data set,” on page 349.
- Macro updates; see Appendix K, “SMF C and assembler header macros,” on page 365.
- Locating information; see “Access to documents” on page xiii and “How to look up message explanations with LookAt” on page xv.

Moved information

This edition includes moved information:

- Client support; see “NFS client translation support” on page 85 and “z/OS NFS client with z/OS NFS server” on page 86.
- Reader’s comment information into the Edition notice information.

Deleted information

This edition deleted information:

- DOS and PC support references; Chapter 10, “Customization,” on page 119, Chapter 12, “Installation-wide exits for the z/OS NFS server,” on page 173, Appendix J, “PCNFSD protocol,” on page 361, and Appendix I, “Retrieval of source code for client enabling commands,” on page 359.

Part 1. Introduction

Chapter 1. Introduction 3
Overview 3
z/OS UNIX files 4
z/OS UNIX advantages. 4
NFS protocol compliance 5
Conventional MVS data sets 5
Mounting MVS data sets onto a client mount point 5
Crossing between file systems–NFS server . . . 6
Creating conventional MVS data sets 6
Serializing and sharing data sets. 6
Server control files 7
Attributes data set 7
Exports data set 7
Mount handle data set 8
Log data set 8
Checklist data set. 8
| Tested clients for the z/OS NFS server. 9
| NFS protocol attributes for the z/OS NFS server . 9
Supported servers for the z/OS NFS client 9
WebNFS support 9
| NFS version 3 and version 4 with TCP/IP protocols 10
| TCP/IP version 6 11
Data transfer and conversion 11
Native ASCII support - NFS V2 and V3 11
User-specified port range support 12
| Dynamic addressing 12

© Copyright IBM Corp. 1991, 2006

Chapter 1. Introduction

This chapter explains the Network File System's client-server relationship and introduces the IBM Network File System (z/OS NFS). When used to access z/OS UNIX System Services (z/OS UNIX) data, which conforms to portable operating system interface (POSIX) standards, it is similar to other UNIX/AIX® Network File Systems.

Overview

A client is a computer or process that requests services on the network. A server is a computer or process that responds to a request for service from a client. A user accesses a service, which allows the use of data or other resources.

Figure 1 illustrates the client-server relationship. The upper right portion of the figure shows the z/OS NFS server. The lower right portion of the figure shows the z/OS NFS client. The left portion of the figure shows various NFS clients and servers which can interact with the z/OS NFS server and client. The center of the figure shows the Transmission Control Protocol/Internet Protocol (TCP/IP) network used to communicate between the clients and servers.

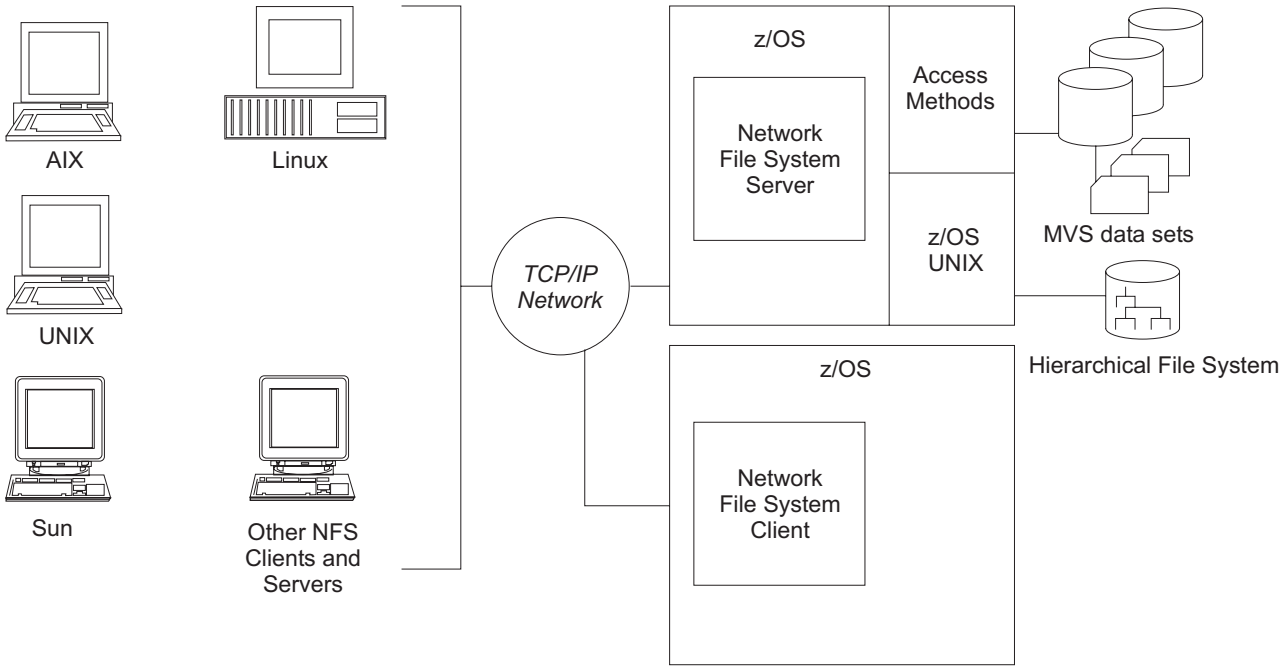

Figure 1. NFS client-server relationship

With the z/OS NFS server, you can remotely access MVS/ESA™ conventional data sets or z/OS UNIX files from workstations, personal computers, and other systems that run client software for the Sun NFS version 2, version 3, and version 4 protocols, and the WebNFS protocols over TCP/IP network.

The z/OS NFS server acts as an intermediary to read, write, create or delete z/OS UNIX files and multiple virtual storage (MVS) data sets that are maintained on a z/OS host system. The remote MVS data sets or z/OS UNIX files are mounted

© Copyright IBM Corp. 1991, 2006

from the host processor to appear as local directories and files on the client system. This server makes the strengths of an z/OS host processor (storage management, high-performance disk storage, security, and centralized data) available to the client platforms.

With the z/OS NFS client you can allow basic sequential access method (BSAM), queued sequential access method (QSAM), virtual storage access method (VSAM), and z/OS UNIX users and applications transparent access to data on systems which support the Sun NFS version 2 protocols and the Sun NFS version 3 protocols. The remote NFS Server can be a z/OS, UNIX, AIX, or other system. The z/OS NFS client is implemented on z/OS UNIX and implements the client portion of the Sun NFS version 2 protocols and the Sun NFS version 3 Protocols.

The NFS uses the communication services provided by TCP/IP, a suite of protocols that includes the remote procedure call (RPC) and External Data Representation (XDR) protocols. RPC allows a program on one machine to start a procedure on another machine, as if the procedure is local. XDR resolves the differences in data representation of different machines.

The NFS, then, can be used for file sharing between platforms and file serving (as a data repository).

If you use NFS as a file server, the z/OS UNIX file system might be a better choice than using conventional MVS data sets, because of its UNIX-based features.

z/OS UNIX files

The NFS server enables the client user remote access to z/OS UNIX files from a client workstation.

z/OS UNIX provides a hierarchical file system (HFS) for z/OS. The HFS file system is similar to a UNIX file system. All z/OS UNIX files reside in a directory, which in turn is a file in a higher level directory. The highest level directory is called the root directory.

When client users mount files from your server system, you use a common HFS prefix to distinguish z/OS UNIX files from conventional MVS data sets. You see z/OS UNIX files in a standard UNIX format on your workstation, but the files are stored on a z/OS host system.

Using the NFS, the client can mount all or part of the z/OS UNIX file system and make it appear as part of your local file system. From there the client user can create, delete, read, write, and treat the host-located files as part of the workstation's own file system. For more information about z/OS UNIX see *z/OS UNIX System Services User's Guide*.

z/OS UNIX advantages

z/OS UNIX file system support provides these advantages over conventional MVS data sets:

- Support for hierarchical directories
- File names up to 255 characters in length
- Path names up to 1023 characters in length
- Mixed case names and special characters, except nulls and slash characters, in file and path names

- UNIX style access permission support
- Group and user ID support at a file level
- Ability to link conventional MVS data sets to a POSIX path name.

NFS protocol compliance

The z/OS Network File Systems provides full NFS protocol compliance for accessing the z/OS UNIX file system.

Conventional MVS data sets

Using NFS, you can access conventional MVS data sets from a client workstation, personal computer, or any client system using software for the NFS protocol.

In MVS, a file is called a data set. The NFS allows client users to mount conventional MVS data sets from their workstations. It presents the information to them in the form of a UNIX (or AIX) file, though the information is actually stored on an MVS-owned DASD.

The files for an operating system are organized into a file system. The UNIX environment use a file system that is a hierarchy of directories. Conventional MVS, in contrast to z/OS UNIX, uses a non-hierarchical file system in which groups of data sets are referred to by specifying a high-level qualifier (HLQ).

The MVS HLQ can include the first (leftmost) qualifier of data sets, or the first and second qualifiers, or the first, second, and third qualifiers, and so on. For example, `SMITH` is the HLQ for the files named `SMITH.TEST.DATA` and `SMITH.PROJ7.SCHED`, while `SMITH.TEST` is the HLQ of `SMITH.TEST.DATA` and `SMITH.TEST.DOCS`.

Mounting MVS data sets onto a client mount point

To access an MVS file system from the client, client users use the **mount** command to create a temporary link (until unmounted) between specific MVS data sets and their UNIX directory (preferably empty) or an unused logical drive on their workstations. The empty UNIX directory or logical drive is called a mount point.

Client users use an MVS HLQ in the **mount** command to specify which MVS data sets to mount at a mount point. The MVS data sets beginning with the specified HLQ appears as files under the mount point.

Client users can also perform a mount using a fully qualified data set name or an alias to a user catalog, but not the catalog name itself. Only cataloged data sets are supported by the z/OS NFS server. Tape data sets and generation data sets are not supported.

Some client platforms support both TCP and the user datagram protocol (UDP). Users can choose either TCP or UDP to access the server. The default protocol option depends on the NFS client platform. For NFS version 4, some platforms do not support UDP.

Note: When directly mounting on a fully qualified data set name, the server must return the mount size as part of getting the attributes for the mount. This can slow down the completion of the **mount** command.

Clients using the NFS version 4 protocol must pass mount requests to the server using a series of **lookup** operations. Client users may still use the mount command, and the client must convert the mount command into the lookup operations.

Crossing between file systems–NFS server

Crossing file systems means the NFS client can also potentially be a server, and remote and local mounted file systems can be freely mixed. This leads to some problems when a client travels down the directory tree of a remote file system and reaches the mount point on the server for another remote file system. Allowing the server to follow the second remote mount would require loop detection, server lookup, and user revalidation. When a client does a lookup on a directory on which the server has mounted a file system, the client sees the underlying directory instead of the mounted directory.

The NFS server does not support crossing file systems in NFS protocol versions 2 and 3. For example, if a server has a file system called /usr and mounts another file system on /usr/src, a client can also mount /usr, but the server will not see the mounted version of /usr/src. A client could perform remote mounts that match the server's mount points to maintain the server's view. In this example, the client would also have to mount /usr/src in addition to /usr, even if the mounts are from the same server.

For NFS protocol version 4, the NFS server does support crossing file systems. This change is due to the mount protocol being obsolete in NFS version 4. Using the example of the previous paragraph, the server would see the mounted version of /usr/src and any objects within that file system.

Creating conventional MVS data sets

Client users can create MVS data sets from a client system using the NFS. The default data set creation attributes specified by the system administrator are used to create MVS data sets, unless the user overrides them. These attributes determine how the MVS data sets are structured and where they are stored. Client users can override the default data set creation and processing attributes for a mount point when issuing the **mount** command. In addition, you can override these attributes at file creation time.

Serializing and sharing data sets

The z/OS NFS server handles data set serialization and sharing differently, depending on the type of data set:

Physical sequential
: The server ensures physical sequential data set read/write integrity by SVC 99 dynamic allocation with exclusive option whenever a physical sequential data set is opened for output. Otherwise, it allocates with share option.

Virtual storage access method (VSAM)
: The server dynamically allocates a VSAM data set with share option and allows the VSAM access method to manage data sharing using the **shareoptions** specified during data set definition.

Partitioned data set extended (PDSE)
: The server dynamically allocates a PDSE data set with share option and allows the PDSE functions to manage the serialization of the PDSE data set and its members.

Partitioned data set (PDS)
: For read and write, the z/OS NFS server issues ENQ SHR on QNAME=SYSDSN and RNAME=dataset_name (through an SVC 99). For write, the server issues an exclusive ENQ against QNAME=SPFEDIT and RNAME=dataset_name.member_name, in addition to the serialization of resources by SVC 99. For all MVS users who are allocating their data set with exclusive status, this provides write protection. It only provides read integrity for ISPF users.

Server control files

These special files are used by the z/OS system administrator to control the z/OS NFS server:

- Attributes data set
- Exports data set
- Mount handle data sets
- Log data sets
- Checklist data set

For information about customizing these control files, see "Configuring the z/OS NFS server" on page 133.

Attributes data set

The attributes data set contains the settings for the z/OS NFS server. There are three types of attributes stored in this data set:

Data set creation attributes
: Used to define the structure of MVS data sets when creating a file (for conventional MVS data sets only).

File processing attributes
: Used to control how files are accessed by the client.

Site attributes
: Used to control z/OS NFS server resources.

The system administrator changes the default settings by editing the attributes data set and restarting the server. Client users can override the data set creation and file processing attributes at the command line. For conventional MVS data sets, the client user can specify the data set creation attributes when mounting, creating, or accessing files. The client user can override the file processing attributes when mounting, creating, or accessing files. However, some file processing attributes can only be overridden on a mount point basis.

Note: Many of the attributes are valid for conventional MVS files only, and not for z/OS UNIX files. "Attributes used for z/OS UNIX System Services file access" on page 87 gives a complete list of attributes that are valid for z/OS UNIX.

Exports data set

The exports data set can control which client users can mount which MVS data sets. The entries in the exports data set specify which MVS high-level qualifiers or HFS directories can be mounted. The system administrator can use this data set to limit mounts to accredited clients only. It also controls which client users can mount all or part of the z/OS UNIX file system, based on the client machine's

specified Internet Protocol (IP) address. To use the exports data set, the **security** site attribute must be set to either **safexp** or **exports** by the MVS system administrator.

Mount handle data set

The z/OS NFS server maintains a list of the active mount points in a pair of files called the mount handle data sets on MVS. The two data sets are used alternately to automatically reestablish the client mount points when the server is started. If the file system is not available, the mount point is not reestablished and the mount failure is recorded in the log data set.

The z/OS NFS server does the cleanup activity during z/OS NFS server shutdown and daily at the cleanup time specified by the **restimeout** site attribute.

During cleanup time, the z/OS NFS server reads the list and checks all mount points against the retention period specified in the restimeout site attribute. If your mount points are idle longer than the retention period specified in the **restimeout** site attribute, they are removed. Only the active mount points are reconnected.

If a mount handle is removed by the cleanup activity, the client user might receive the "Stale NFS File Handle" message or some other appropriate message. If so, all the client user needs to do is unmount the stale mount point and mount it again.

Log data set

The log data sets store the messages for the z/OS NFS start-up procedures. This log can be used to identify the user's correctable errors or the user's problem errors. There are two logs that this information is stored in; the primary log and the secondary log. The primary log is used at start-up until it is filled and then overflows into the secondary log. When the secondary log is full, the primary log will then be overwritten with new error messages. The number of log records is dependent on the number of transactions that the server can handle.

The z/OS NFS server also records messages and diagnostic information in a z/OS component trace buffer, if one is specified. Component trace buffers can be used in addition to or instead of the log data sets. Using a component trace buffer can provide performance improvements over the log data sets. For details, see "Using z/OS component tracing" on page 229.

Checklist data set

The checklist data set contains entries for files or directories that are to be exempt from System Authorization Facility (SAF) checking even though **saf** or **safexp** is specified as the security options. This file is only used when **saf** or **safexp** is specified for the particular data type and the **checklist** option is specified as a site attribute. The entries specified here must match a subsequent mount point, or be the parent of a subsequent mount point, to allow SAF checking to be bypassed for everything underneath that mount point. If the entry does not match a subsequent mount point and it is not a parent of any subsequent mount point then it has no effect.

Tested clients for the z/OS NFS server

Tested clients for the z/OS NFS server, using the NFS version 4 protocol, are:

- IBM RS/6000® AIX version 5.3
- Sun Solaris version 10
- Enterprise Linux 4
- Windows 2000/XP with Hummingbird Maestro 9 and Maestro 10
- Other client platforms should work as well since NFS version 4 is an industry standard protocol, but have not been tested by IBM.

Older versions of the above clients are still supported under the NFS version 2 and version 3 protocols, but all have not been tested by IBM.

NFS protocol attributes for the z/OS NFS server

The NFS protocol defines file attributes that NFS clients can read and set on NFS servers. In the NFS version 4 protocol, some file attributes are mandatory and others are recommended for servers to support. For a list of the NFS version 4 file attributes that the z/OS NFS server supports, see Appendix C, "NFS attributes," on page 321.

Supported servers for the z/OS NFS client

The z/OS NFS client supports all servers that implement the server portion of the Sun NFS Version 2 and Version 3 protocols. The z/OS NFS client does not support NFS version 4.

A **mount** parameter **vers(x)**, where **x** is either 2 or 3 is provided to make the z/OS NFS client communicate with the server at the specified protocol level. The z/OS NFS client also communicates at the highest protocol level that is supported by the server if no level is specified.

- If **no version** is specified and if the server:
  - Only supports the NFS version 2 protocol, then the z/OS NFS client will use the NFS version 2 protocol to communicate
  - Supports both the NFS version 2 and 3 protocols, then the z/OS NFS client will use the NFS version 3 protocol to communicate
  - Supports the NFS version 2, 3 and 4 protocols, then the z/OS NFS client will use the NFS version 3 protocol to communicate.
- If **vers(2)** is specified, then use NFS version 2 protocol to communicate with the server.
- If **vers(3)** is specified, then use NFS version 3 protocol to communicate with the server. z/OS NFS client fails the mount command if the server does not support NFS version 3 protocol.

WebNFS support

The z/OS NFS server supports the WebNFS protocol. WebNFS specification extends the semantics of NFS versions 2, 3 and 4 protocols to allow clients to obtain file handles without the mount protocols. The z/OS NFS server supports the public filehandle and multi-component lookup features as well as other additional requirements as described in RFC 2055. A keyword, **public**, is added for the system administrator to specify the public paths that the public file handle can access. A public path for conventional MVS data and a public path for HFS data

can both be specified. When a lookup request comes in from an NFS client and an absolute path name is specified, it will be matched with the public paths to determine which public path it is trying to reference. If a relative path is specified and both HFS and MVS public paths are defined then the lookup request will be processed relative to the HFS public path.

The following are restrictions for the WebNFS support provided by the z/OS NFS server in this release.

Export Spanning Pathnames
Lookup requests, that reference files or directories outside of the exported public path, will result in an error condition.

Symbolic Links
A symbolic link embedded in a multi-component pathname lookup request will result in an error condition. However, if the final component is a symbolic link, the server will return the filehandle of the symbolic link and let the client evaluate it. External links, which are special cases of symbolic links, will be handled similarly.

Native Path
Only canonical pathnames will be supported.

Canonical path
A canonical path is a hierarchically-related, slash-separated sequence of components, in the form: <directory>/<directory>/../<name>.

NFS version 3 and version 4 with TCP/IP protocols

Information for NFS version 3 and version 4 protocols with **proto=tcp** can be found on the mount man page on a UNIX client. The NFS client automatically selects the option unless the end-user overrides the option. For example, you can enter this command:

```
unix$ mount -o vers=2,proto=udp mvshost1:smith /mnt
```

The above example shows the preference of NFS version 2 with **udp** protocol, even though the client platform can handle the NFS version 3 and **tcp** protocol.

Users can issue the `rpcinfo -p <hostname>` to show all the RPC programs available on the server. Table 2 shows an example.

```
$ rpcinfo -p mvshost1
```

Table 2. View of NFS server capability

program	vers	proto	port	service
100000	2	udp	111	portmapper
100000	2	tcp	111	portmapper
100003	2	udp	2049	nfs
100003	3	udp	2049	nfs
100059	2	udp	2012	
100044	1	udp	2013	
100005	1	udp	2014	mountd
100005	3	udp	2015	mountd
150001	1	udp	2016	pcnfsd
150001	2	udp	2017	pcnfsd
100044	1	tcp	2010	
100059	2	tcp	2011	
100003	2	tcp	2049	nfs

Table 2. View of NFS server capability (continued)

program	vers	proto	port	service
100003	3	tcp	2049	nfs

TCP/IP version 6

TCP/IP version 6 expands the range of addresses that are available for internet communications. TCP/IP version 6 extends address sizes from a 32-bit value to a 128-bit value, vastly expanding the number of globally unique addresses that can be assigned. z/OS V1R7 NFS allows both the z/OS NFS client and the z/OS NFS server to support the longer addresses of TCP/IP version 6, as well as the 32-bit addresses of TCP/IP version 4 and below. Your network infrastructure must be enabled to use TCP/IP version 6; if the network does not support IP version 6, z/OS NFS falls back to use IP version 4.

Use of TCP/IP version 6 depends on the NFS protocol version level. The z/OS NFS server can use TCP/IP version 6 for all NFS functions when communicating with the NFS version 4 protocol. With the NFS version 2 or version 3 protocols, the z/OS NFS client is restricted to TCP/IP version 4, which supports the portmapper program that clients must use to identify communication ports. For NFS version 4 the NFS client doesn't use the portmapper, so the client can communicate with the NFS server using NFS version 4 and TCP/IP version 6. However, for the NFS client to use the client enabling utilities, such as mvslogin, mvslogout, and showattr, the client reverts to TCP/IP version 4 and attempts to identify the ports via the server portmapper. Once the ports are identified, utility communication with the NFS server is done using TCP/IP version 4.

SMF records for the z/OS NFS server report client IP addresses for both TCP/IP version 4 and version 6, with separate address fields for each.

Data transfer and conversion

With the NFS version 4 protocol, text data and metadata are transferred between the server and client in the UTF-8 data format. With NFS V4, ASCII text is not transferred directly.

Native ASCII support - NFS V2 and V3

With the NFS version 2 and 3 protocols, the z/OS NFS client and server support applications running on z/OS V1R2 (and higher) in a Native ASCII environment. Applications can operate on files in either EBCDIC or ASCII format as well as other data formats defined with a coded character set identifier (CCSID). Native ASCII support is provided with a mechanism called file tagging where the file is defined with a tag to identify the CCSID to use for data conversion. File tagging is defined in the appropriate z/OS UNIX System Services documents. The z/OS NFS client and server provide the necessary support to provide data conversion between different CCSIDs specified for the client and server. The z/OS NFS client **cln_ccsid** and **srv_ccsid** parameters will also be supported by the z/OS NFS server to identify the CCSID to be used in the data conversion. See "Processing attributes syntax" on page 93 for more information about the **cln_ccsid** and **srv_ccsid** parameters.

User-specified port range support

The z/OS NFS server supports a user-specified range of ports. The z/OS NFS server allows users to specify port assignments for services **mountd**, **mvsmount**, **pcnfsd**, and **showattr**. Additional ports are also required by the server for locking functions. The port assignments for these services can be any port number (except for reserved port 2049 for the NFS program) but must be a contiguous port range for the z/OS NFS server to identify them. The user specified range of ports provides a flexible port range to accommodate programs such as a firewall that supports a range of ports for security purposes.

Users wanting a user-specified port range setup must change the /etc/rpc file for the z/OS NFS client and the /etc/services and tcpip.profile files for the z/OS NFS server. For more information, see "Setting up a user specified port range" on page 143.

Dynamic addressing

Before z/OS V1R7, the z/OS NFS client and server were based on the static IP address model to handle all communications with other systems. However, many systems have migrated from the use of static IP addresses to the dynamic host configuration protocol (DHCP). Now, the z/OS NFS server accepts dynamic NFS client IP address changes and properly understands the source of the communication even if the sender's IP address has changed. Since not all customers' environments use dynamic IP addresses, NFS server site attributes have been added to specify whether the NFS server should use the dynamic IP algorithm (**dhcp**) or the current static IP algorithm (**nodhcp**). The default is **nodhcp**, to use the static IP algorithm.

To use dynamic IP addressing, the client must:

- Have a constant host name that the NFS server can identify it by.
- Dynamically update the authentication DNS (dynamic name server) with new IP addresses whenever they change.
- Maintain the TTL (time to live) value that the authentication DNS server specifies to any caching DNS server, based on the frequency with which system IP addresses might change.

For more information, see "Using dynamic client IP addressing" on page 147.

The z/OS NFS Server continues to have a static IP address, based on the standard industry practice of assigning static IP addresses to servers.

Part 2. NFS User's Guide

Chapter 2. Creating conventional MVS data sets 15
Overriding data set creation attributes 15
Preparing to create an MVS file. 15
Naming MVS files 16
Restrictions using alias names for MVS files 16
Creating physical sequential files 16
Creating direct access files 17
Creating PDSs and PDSEs 17
Creating a PDS or PDSE - mkdir dsntype(pds), dsntype(library) 17
Removing a PDS or PDSE - rm, rmdir 18
Accessing PDS or PDSE members 18
Updating or extending a PDS or PDSE member 19
Timing out while writing a PDS or PDSE member 19
Wildcard copy to a PDS or PDSE 19
Limitations of a PDS 19
Concurrent writes to a PDSE 19
Creating VSAM files 19
Exploiting SAM striped files. 20
Exploiting large format data sets 21

Chapter 3. Using conventional MVS data sets . . 23
Special MVS considerations 23
Selection of an MVS data storage format . . . 23
File size determination and time stamps. . . . 24
Ownership and permissions. 24
NFS version 2 and version 3 statelessness . . . 24
NFS version 4 state. 25
File reading and writing 26
Random access to files. 26
Cached data writing 27
Case sensitivity–maplower, nomaplower. . . . 27
Selection of text or binary processing modes–text, binary 27
Binary processing mode 27
Text processing mode 28
Number representation 28
Access to migrated files–retrieve, noretrieve; wait, nowait 28
Access to migrated system-managed data sets . . 29
File handle refresh 29
Mapping between the workstation and MVS file systems. 30
File name extension mapping 30
Mounting of MVS data sets onto a client mount point 31
Mount examples. 31
Variants of the mount command 32
Use of a PDS or PDSE as a directory 33
Use of multiple mount points 33
Data set serialization and sharing 34
NFS protocol 34
NFS file system attributes 34

Chapter 4. Using z/OS UNIX System Services files. 35
z/OS UNIX file system 35
POSIX compatibility 36
NFS protocol 36
Attributes specific to z/OS UNIX System Services 36
Synchronous write to a z/OS UNIX file for NFS version 2 protocol 37
Synchronous write to a z/OS UNIX file for NFS version 3 or 4 protocol. 37
Authorization checking when writing to a z/OS UNIX file 37
HFS site attribute 37
Protecting your z/OS UNIX System Services files . 38
Accessing z/OS UNIX System Services files from a client 38
Mount examples. 38
z/OS UNIX data transfer and conversion 39
Data transfer under the NFS version 4 protocol 39
Text or binary processing - NFS version 2 and 3 protocols 39
Linking an MVS data set to a z/OS UNIX file system 40
Creating an external link 40
Displaying the contents on an external link. . . 41
Deleting an external link 41
UNIX look and feel. 41
NFS file system attributes 42

Chapter 5. Locking and access control 43
Using Network Lock Manager (NLM) in NFS V2 and V3 43
Monitored locks 43
Non-monitored locks 43
Specifying a grace period for reclaiming locks . . 44
Listing locks held for a file 44
Releasing locks held for a file 44
Using Network Status Monitor (NSM) in NFS V2 and V3 44

Chapter 6. Commands and examples for AIX and UNIX clients 47
Using commands on AIX. 47
Quick reference of AIX and UNIX commands . . . 49
Accessing z/OS UNIX System Services and conventional MVS files 50
Mvslogin command examples 51
"Permission denied" message 51
Mount command examples 52
Overriding default attributes 52
Using mount to override server default attributes 53
Getting authorization to access files 53
Saving of mount points 53
Automatic timed logout - logout attribute . . 54
Displaying default and mount point attributes - showattr 54

© Copyright IBM Corp. 1991, 2006

Unmounting and logging out of z/OS 57
Disconnecting your mount point - umount . . . 57
Ending your z/OS session - mvslogout 58

Chapter 7. Commands and examples for z/OS NFS clients 59
Using commands on DFSMS/MVS 59
Accessing z/OS 63
Mvslogin command examples 64
"Permission denied" message 64
Mount command syntax and examples 65
Data conversion 66
BSAM, QSAM, and VSAM ESDS access to remote files 67
Mount command examples 68
Getting authorization to access files 70
Saving of mount points 70
Automatic timed logout - logout attribute . . 70
Unmount command syntax and examples . . . 71
Disconnecting your mount point - unmount 71
Displaying client statistical information–nfsstat 72
Displaying server mount information–showmount 75
Displaying default and mount point attributes–showattr 76
Ending your z/OS session - mvslogout 77

Chapter 8. Initialization attributes for the z/OS NFS client 79
Client attribute syntax 79
Datacaching attribute 84
Mount processing parameters and installation parameters 84
NFS client translation support 85
z/OS NFS client with z/OS NFS server 86

Chapter 9. Initialization attributes for the z/OS NFS server 87
Attributes used for z/OS UNIX System Services file access 87
Multipliers 87
Duplicate attributes. 88
Data set creation attributes syntax 88
Processing attributes syntax 93
Timeout attributes 99
Retrieve attributes 99
Mapped keyword processing attribute 100
Native ASCII processing attributes 100
Considerations for Native ASCII environment support 100
NFS servers with non-z/OS based NFS clients 101
Site attributes syntax 104

Chapter 2. Creating conventional MVS data sets

This chapter explains how to create the various types of data sets (files) that are supported by the z/OS NFS server.

The examples shown are for the AIX RS/6000 platform. Any examples for other platforms are so indicated.

Overriding data set creation attributes

When you create an MVS file, default file creation attributes are applied, unless you override them. The attributes are passed to the z/OS host.

Data set creation attributes are controlled in the following ways, in increasing order of priority.

- Default server data set creation attributes
- Default installation data set creation attributes, specified by the system administrator in the attributes data set
- DFSMS data class attributes
- Data set creation attributes specified in the **mount** (or **nfs link**) command
- Data set creation attributes specified in the **mkdir**, **vi** (edit), or **cp** (copy) commands (highest priority)

Note: The z/OS NFS server does not support the following data class attributes.

- Retention period/Expiration date
- Number of volumes
- VSAM imbed index option
- VSAM replicate index option
- CI size of data component
- Percentage of CI or CA free space

Preparing to create an MVS file

When creating an MVS file, you should know whether to process the file in text or binary mode (see "Selection of text or binary processing modes–text, binary" on page 27) and what type of file to create.

The z/OS NFS server supports the following types of files.

- Physical sequential (PS) data sets, including basic format and extended format data sets but excluding compressed format data sets.
- Direct access (DA)
- Partitioned data sets (PDS)
- Partitioned data sets extended (PDSE)
- VSAM KSDS
- VSAM ESDS
- VSAM RRDS

Keyed access to files and GDG are not supported.

© Copyright IBM Corp. 1991, 2006

Naming MVS files

The NFS uses the comma (,) as a delimiter to a list of file attributes. Do not use a comma as a special character in file name. For example, you can enter this command:

```
$ vi "/u/smith/new,text"
```

This indicates to NFS that a file called *new* is being edited in the attribute text mode, not file `new,text`.

When naming conventional MVS files, you must follow the MVS file naming conventions, as described in *z/OS DFSMS Using Data Sets*.

For information about the z/OS UNIX System Services naming conventions, see Chapter 4, "Using z/OS UNIX System Services files," on page 35.

An MVS file name (or data set name) can consist of one or several simple names joined so that each represents a level of qualification. For example, the MVS file name `DEPT58.SMITH.DATA3` is composed of three qualifiers.

The following characteristics apply to the MVS file name.

- Each qualifier consists of 1 to 8 alphanumeric characters, national characters (@, #, $), or a hyphen (-)
- Each qualifier must start with an alphabetical or national character
- The period (.) separates simple names from each other
- Including all simple names and periods, the length of the MVS file name must not exceed 44 characters
- PDS and PDSE member names can be up to 8 characters long

For information about the MVS file system, see "Mapping between the workstation and MVS file systems" on page 30.

Restrictions using alias names for MVS files

For non-VSAM files, alias names can be used interchangeably with the true file name except on remove (**rm** and **rmdir**) and rename (**mv**) requests. Renaming or removing an alias name results in an I/O error. This is due to an MVS restriction.

If the true name of an MVS file is renamed or removed and alias names have been defined for the file, MVS deletes the alias names during execution of the rename or remove request.

Creating physical sequential files

When creating a physical sequential (PS) file, specify the **dsorg(ps)** attribute (if it is not the default already) with the **mount** command or a file creation command, such as the **vi** UNIX (or AIX) command.

1. Create a local directory on your client to be used as a mount point. For example (with UNIX), enter this command:

   ```
   $ mkdir /u/smith/mnt
   ```

2. Mount the MVS file system. For example, suppose your host is `mvshost1`, and you want to issue a mount on the high-level qualifier `smith`. You can enter this command:

   ```
   # mount mvshost1:"smith,dsorg(ps)"  /u/smith/mnt
   ```

If you do not specify any other attributes, the MVS site defaults are used. You can use the **showattr** command to display the site defaults.

3. You can use the **vi** UNIX command to create the `new` file.

   ```
   $ vi /u/smith/mnt/new
   ```

   When you save the file using **vi**, you have just created a new MVS PS file named `SMITH.NEW`.

You can get the same results by specifying **dsorg**(*ps*) in the file creation command rather than in the **mount** command.

```
# mount mvshost1:smith /u/smith/mnt
$ vi "/u/smith/mnt/new,dsorg(ps)"
```

Creating direct access files

When creating a direct access (DA) file, specify the **dsorg(da)** attribute (if it is not the default already) with the **mount** command or a file creation command (such as the **vi** UNIX command).

1. Create a local directory on your client to be used as a mount point. For example (with UNIX), enter this command:

   ```
   $ mkdir /u/smith/mnt
   ```

2. Mount the MVS file system. For example, suppose your host is `mvshost1`, and you want to issue a mount on the high-level qualifier `smith`. You can enter this command:

   ```
   # mount mvshost1:"smith,dsorg(da)"  /u/smith/mnt
   ```

   If you do not specify any other attributes, the multiple virtual system (MVS) site defaults are used. You can use the **showattr** command to display the site defaults.

3. Next, you can use the **vi** UNIX command to actually create the `new` file.

   ```
   $ vi /u/smith/mnt/new
   ```

   You have just created a new MVS DA file named `SMITH.NEW`.

You can get the same results by specifying **dsorg(da)** in the file creation command, rather than in the **mount** command.

```
# mount mvshost1:smith /u/smith/mnt
$ vi "/u/smith/mnt/new,dsorg(da)"
```

Creating PDSs and PDSEs

Partitioned data sets (PDS) and partitioned data sets extended (PDSE) can be used as directories, and their members are files within those directories. An illustration of the use of PDSs to act as directories is shown in Figure 2 on page 33. For general information on PDSs and PDSEs, see *z/OS DFSMS Using Data Sets*.

You cannot create new directories within a PDS or PDSE, due to the nature of these data structures.

Creating a PDS or PDSE - mkdir dsntype(pds), dsntype(library)

To create a PDS or PDSE, perform the following steps:

1. Create a local directory on your client to be used as a mount point. For example (with UNIX), enter this command:

   ```
   $ mkdir /u/smith/mnt
   ```

2. Mount the MVS file system (accessing files that begin with the high-level qualifier of smith).

   ```
   # mount mvshost1:"smith,mgmtclas(normal)"
   /u/smith/mnt
   ```

3. If creating a PDSE, use the **mkdir** (make directory) UNIX command, specifying the **dsntype(library)** attribute to create a PDSE named smith.datalib.

   ```
   $ mkdir /u/smith/mnt/"datalib,dsntype(library)"
   ```

   If creating a PDS, use the mkdir (make directory) UNIX command, specifying the following dsntype(pds) attribute.

   ```
   $ mkdir /u/smith/mnt/"datalib,dsntype(pds),dir(20)"
   ```

   Omitting dsntype(pds): You can omit specifying the **dsntype(pds)** attribute if **pds** has been specified for the **dsntype** attribute either in your site attribute data set or in your mount point.

4. You can use the **vi** UNIX command to create a PDS or PDSE member named smith.datalib(member1).

   ```
   $ vi "/u/smith/mnt/datalib/member1,text"
   ```

   Type your text, save it, and quit.

You have now created a PDS or PDSE member, which is processed in text processing mode. You can use the **cat** UNIX command to view the contents of your PDS or PDSE member.

Note: z/OS NFS server supports a maximum of 14,562 members in a PDS or PDSE data set. When a NFS read-directory request on a PDS or PDSE is processed, the z/OS NFS server will return up to 14,562 member names. Other requests, such as read and write, to individual members are not affected.

Removing a PDS or PDSE - rm, rmdir

To remove a PDS or PDSE, first make sure that the PDS or PDSE is empty. You can delete all members under the directory using the **rm** UNIX command. Then use the **rmdir** (remove directory) UNIX command. This example removes the datalib directory, and confirms its removal by a failed try to query it (**ls** is the UNIX list files command).

```
$ ls -F /u/smith/mnt/datalib
data1* data2* data3*
$ rm /u/smith/mnt/datalib/*
$ rmdir /u/smith/mnt/datalib
$ ls -F /u/smith/mnt/datalib
/u/smith/mnt/datalib not found
```

Accessing PDS or PDSE members

There is more than one way to mount and access PDS and PDSE members. For example, you can display the existing PDS member smith.source(bigblue) by entering either of these command sequences.

```
$ mkdir /mnt
# mount hostname:"smith.source,text" /mnt
$ cat /mnt/bigblue
```

Or

```
$ mkdir /mnt
# mount hostname:"smith,text" /mnt
$ cat /mnt/source/bigblue
```

These two approaches are equivalent.

Updating or extending a PDS or PDSE member

The z/OS NFS server does not support updating or extending a PDS or PDSE member directly. To update or extend a PDS or PDSE member, a client program must follow these steps.

Step 1. Copy the file to the client machine

Step 2. Update or extend the copied version on the local system

Step 3. Truncate the original MVS file to zero size by sending a `SETATTR` request with zero file size

Step 4. Copy the updated version on the local host to MVS by writing request

Some client editors follow the above steps, for example, the AIX and UNIX **vi** editor. Other editors do not follow the above steps, for example, the z/OS UNIX OEDIT editor. In the latter case the user must save the updated version into a new file.

Timing out while writing a PDS or PDSE member

If you are writing to a PDS or PDSE member and a timeout occurs, the timeout causes the member to close. The remaining write requests appear to append to a PDS or PDSE member. The write request does not complete successfully and causes an I/O error. To avoid timing out, increase the time on the timeout setting.

Wildcard copy to a PDS or PDSE

To ensure that a wildcard copy of a PDS or PDSE is completed successfully, the PDS or PDSE member must be closed and dequeued (if necessary). For example, using the statement below, a wildcard copy will fail if any member inside of `smith.datalib` is open.

```
$ cp smith.*  /u/smith/mnt/datalib
```

Limitations of a PDS

The PDS support in NFS adheres to the conventions used in MVS. For example, you cannot have more than one member of a PDS open for output at a time. If you try to create, remove, rename, or write a member of a PDS while another member is open for output, you get a “Permission denied” message.

A PDS member stays open for the timeout period specified in the appropriate timeout processing attribute, or until you try to create or write to another member.

Concurrent writes to a PDSE

NFS supports concurrent writes to a PDSE. If you are writing to one member of a PDSE, another client can write to any other member in the same PDSE. You can also use ISPF to edit a PDSE member while another NFS client is writing to a different member of the same PDSE.

Creating VSAM files

The NFS supports three types of VSAM files: key-sequenced (KSDS), entry-sequenced (ESDS), and relative record (RRDS). However, keyed access and relative-number access to the files are not supported.

If you plan to update a VSAM data set (for example, with the **vi** editor or with the **cp** copy command), the data set must have been defined with the **reuse** option. Trying to write back a VSAM data set that was not defined as reusable results in an "I/O error", "failure to open", or similar error message. If you create a VSAM file using the NFS, the reuse option is used by the server.

For more information on VSAM files, see *z/OS DFSMS Using Data Sets*.

In the following example, the attributes indicate that:

Spanned records are allowed

Organization is key-sequenced

Keys are 8 bytes long and start in position 0 of each record

Average record size is 1024

Maximum record size is 4096

Space is allocated for 50 records with a secondary allocation of 10

Cross-region and cross-system share options are provided

The file is to be created on a volume named D80CAT

```
$ cp ksds.old "ksds.new2,spanned,dsorg(indexed),keys(8,0),
   recordsize(1K,4K),space(50,10),shareoptions(1,3),
   vol(D80CAT)"
```

In the following example for creating a VSAM ESDS file, the attributes indicate that:

Spanned records are allowed

Organization is entry-sequenced

Average record size is 1024

Maximum record size is 4096

Space is allocated for 50 records with a secondary allocation of 10

Cross-region and cross-system share options are provided

The file is to be created on a volume named D80CAT

```
$ cp esds.old "esds.new3,spanned,dsorg(nonindexed),
   recordsize(1K,4K),space(50,10),shareoptions(1,3),
   vol(D80CAT)"
```

In the following example, the attributes indicate that:

- Spanned records are not allowed
- Organization is relative record, numbered in ascending order
- Average record size is 1024
- Maximum record size is 1024
- Space is allocated for 50 records with a secondary allocation of 10
- Cross-region and cross-system share options are provided
- The file is to be created on a volume named D80CAT

```
$ cp rrds.old "rrds.new4,nonspanned,dsorg(numbered),
   recordsize(1K,1K),space(50,10),shareoptions(1,3),
   vol(D80CAT)"
```

Exploiting SAM striped files

With SAM striping, data I/O is done in parallel to improve performance. For a file with 16 stripes, data is simultaneously processed on the first track of the allocated space in each of the 16 volumes. This allows for quick access to all the information.

The z/OS NFS server can support data set striping through the use of data class and storage class attributes that define extended format data sets. The z/OS NFS server can exploit the performance of extended format data sets by reading multiple blocks at a time when reading ahead.

For more information about striped files, see *z/OS DFSMS Using Data Sets*.

Exploiting large format data sets

Large format data sets are a type of physical sequential data set, other than extended format data sets, which can grow beyond a size limit of 65 535 tracks on each volume. That size limit applies to conventional (basic format) sequential data sets. Large format data sets can exploit the increased storage capacity of most hardware storage devices, and reduce the need for very large data sets to span multiple volumes.

To create a large format data set with z/OS NFS, specify a data class which has a dsntype value of *large*. The dsorg value must be ps, psu, or omitted.

For more information about large format data sets, see *z/OS DFSMS Using Data Sets*.

Chapter 3. Using conventional MVS data sets

This chapter explains what you need to know to use conventional MVS data sets on a client workstation. This chapter discusses the following topics:

- Special MVS considerations
- Reading and writing MVS data sets
- Accessing MVS data sets
- Mapping between the workstation and MVS file systems.

In MVS, a file is called a data set, and the two terms are used interchangeably in this book. The z/OS NFS server presents information to you in the form of a UNIX (or AIX) file, though the information is actually stored on MVS-owned DASD in the form of an MVS data set.

Special MVS considerations

In addition to mapping between the workstation and the MVS file systems, the z/OS NFS server might be different from non-MVS NFS servers in other ways, including these differences:

- Selection of an MVS data storage format
- File size determination and time stamps
- Ownership and permissions
- State
- File reading and writing
  - Random access to files
  - Cached data writing
- Case sensitivity–**maplower**, **nomaplower**
- Selection of text or binary processing modes–**text**, **binary**
  - Binary processing mode
  - Text processing mode
- Number representation
- Access to migrated files–**retrieve**, **noretrieve**; **wait**, **nowait**
- Access to system-managed migrated data sets
- File handle refresh
- File name extension mapping

Selection of an MVS data storage format

The files you create with the z/OS NFS server are contained in MVS data sets. These MVS data sets are record-oriented and can be sequential, direct, VSAM, partitioned, and so forth. These MVS data sets are variable or fixed in record length. UNIX files, however, are byte-oriented and typically written or read at certain offsets in the file.

You can map non-MVS files to most types of MVS data set organizations. However, how the time stamps and file size value are handled depends on the type of MVS data set used, and the file size processing can affect performance. See Appendix A, “File size value for MVS data sets,” on page 313.

© Copyright IBM Corp. 1991, 2006

Direct reads with record format **recfm(fbs)** or **recfm(f)** can be fast. In some cases, the z/OS NFS server can determine the physical block addresses from the record offsets. The MVS sequential file organization with **recfm(fbs)** or **recfm(f)** on DASD allows for efficient updating or reading at any offset in the file. Other supported MVS access methods (for example, VSAM) can be used if required by a given application but, in general, the sequential file organization is the best choice for files that are used mainly by UNIX clients.

File size determination and time stamps

How the z/OS NFS server handles the file size value and time stamps depends on the type of MVS data set used and the attributes used to access the data set. See Appendix A, "File size value for MVS data sets," on page 313 and Appendix B, "Time stamps for MVS data sets," on page 317.

Ownership and permissions

The UNIX UID and GID file attributes are reset to their default state after a restart of the z/OS NFS server or an unmount of the file system. In some cases, this requires that a superuser on the client workstation reissue a **chmod** command to reset the UID and GID. This command can be included in the same script used to mount the file system.

The permissions checking done by RACF®, a component of the Security Server for z/OS, or an equivalent security package, is transparent to you. Access to a data set is granted, provided that the server's exports list, the MVS security subsystem, and the customized installation security exit allow access to the data set. Which of these security systems are active depends on the security settings used at your installation. The UNIX file modes or permission bits are ignored by z/OS NFS server and authorization is done with the RACF or equivalent security package.

UNIX's UID, GID, and MODE attributes are not used by the z/OS NFS server for checking user access to conventional MVS data sets (see previous paragraph). UNIX's CHOWN, CHGRP, and CHMOD commands do not update z/OS RACF security policies and will not alter access to files. Do not use returned values to determine access rights. The z/OS NFS server supports the setting and obtaining of these attributes to minimize impact to client applications. For performance, validation of passed values is limited to the following checks for proper operations.

CHOWN
Request is failed with EPERM if changing to a value other than yourself. Change to yourself is allowed for mount support.

CHGRP
No checking.

CHMOD
No checking, new value ignored, existing z/OS NFS server value is left unchanged, "success" is returned to the client.

NFS version 2 and version 3 statelessness

Under the NFS version 2 and version 3 protocols, the z/OS NFS server strives to be as *stateless* as possible; that is, it tries to work correctly without maintaining any state information about any of its clients. However, a failure of the server causes cached writes to be lost, some attributes to be reset, and file handles to become stale, or not valid.

With the NFS version 4 protocol, the z/OS NFS server maintains state information for some client operations, to prevent such loses.

NFS version 4 state

The NFS version 4 protocol introduces *state information* that allows clients and servers to keep track of certain resources.

NFS version 4 uses a value of *clientid* or *stateid* to represent the current state (instance) of client-held resources such as locks, opens, and host restarts. The client and server pass this state information between them on certain operations, allowing both to agree on the current instance of resources held by the client.

NFS version 4 includes new states for the following:

1. Client/Server restart instance
2. Open Share/Deny instance
3. Byte Range Locks instance
4. Client Delegation instance.

The NFS version 4 state that is passed between the client and server represents a single instance in a dynamically changing environment; it is incremented when a state is changed within a group of held resources (restart, open, or lock). Once state is established on the server, the client returns what it believes is the current state. The server compares the client state to the server state to detect stale and out of order requests.

The client uses the **setclientid** operation to notify the server of its intention to use a particular client identifier for subsequent requests that entail creating lock, share reservation, and delegation state on the server. Upon successful completion the server returns a shorthand *clientid* which, if confirmed via a separate step, will be used in subsequent file locking and file open requests. Confirmation of the *clientid* must be done using the **setclientid_confirm** operation to return the *clientid* and *setclientid_confirm* values, as verifiers, to the server.

NFS version 2 and version 3 used the Network Status Monitor (NSM) protocol to determine if resources such as file open share or byte range locks were still in use by a remote client. NFS version 4 no longer uses NSM to communicate a client or server restart. NFS version 4 instead uses a current state on both the client and server, where the state is established and passed in subsequent NFS version 4 operations.

In NFS versions 2 and 3, a client or server issues an NSM **sm_notify** RPC procedure to notify the remote host of a restart. Server resources such as an exclusive byte range lock on a file might remain held until explicitly released by the client. If a client that holds a server resource is removed from the network for a long period without the server being notified, the server resource would be unavailable to other clients until timed out by the server.

NFS version 4 provides a protocol for the client to establish or reestablish state, and associates ownership of subsequent server *stateful* operations to previously established states. To resolve the absent client problem mentioned above, the NFS version 4 client must routinely refresh the state within the server-specified lease time. Upon lease time-out, the server may release resources for the client and make them available to other applications.

- A client obtains the server-specified lease time-out attribute by issuing a **getattr** operation. **getattr** is not a stateful operation, thus it does not require prior state to be established. A **getattr** operation may precede a **setclientid** or **setclientid_confirm** operation.
- Refer to the NFS server's **leasetime** site attribute for setting and tuning lease time periods.

File reading and writing

After the z/OS NFS server is started and you have mounted the MVS data set, you can use regular data access or creation commands from your workstation to access files that reside on MVS.

For example, suppose you accessed an MVS file named `prefix.file3` mounted on the local directory /mnt. This is how you could use the UNIX **cat** command (or a similar command) to display the file:

```
$ cat /mnt/file3
```

Suppose you accessed an MVS file named `prefix.file12` mounted on the local directory /mnt. This is how you could use the UNIX **vi** command (or a similar command) to edit the file.

```
$ vi /mnt/file12
```

Writing a file on MVS is straightforward. If the file already exists, the file's existing attributes are used; they are not modified during the write operation. For the priorities of attributes, see "Overriding data set creation attributes" on page 15.

Random access to files

If your application accesses the files at random offsets, there is a performance implication.

In the UNIX environment, a file is represented as a byte stream. That byte stream is accessible for reading and writing at any byte offset for any byte length. In the MVS environment, a file is represented as a collection of records. The record, rather than a single byte, is the smallest object that can be processed. Therefore, the z/OS NFS server has to convert the byte stream operations from NFS clients into standard access method operations on MVS.

To convert byte stream operations to MVS access method operations, the server has to determine which record in the MVS file contains the offset specified in the NFS read or write request. To determine this, the server reads, mapping byte offsets to records, from the last known location in the file until the record containing the requested byte offset is located.

For example, suppose a file on MVS contains 10,000 variable-length records with a maximum length of 80 bytes for any record. Suppose the first NFS request received tells the server to read 4,000 bytes starting at offset 10,000 bytes. Because the file has not been opened yet, the server would open the file and start reading at the first record, searching for the record that contains offset 10,000. Once it found the record, the server would process the request, which might involve reading more records to find enough bytes to satisfy the request.

Another complication involved in mapping byte offsets to records is the processing defined by the user to apply to a file. For example, if you specify text mode processing with line terminators, the perceived offset into a file from a given client changes.

Cached data writing

The z/OS NFS might cache writes if out-of-sequence data packets are received or if a block of data is partially filled. If the Network File System is processing in the binary data mode, the writes will remain cached until one of the following occurs: the timeout for the request has been reached or the number of cached packets exceeds the number specified in **cachwindow()**. If the NFS is processing text data, the writes remain cached until a timeout occurs. The missing data is padded with binary zeroes and record delimiters so that cached writes for text processing are written in the MVS data set on DASD at the location specified in each cached data packet. In the case of cached data packets for binary processing, only binary zeroes will be used to pad the missing data written at the specified location on DASD. See Table 3 below.

Attention: For the NFS version 3 commit procedure, the z/OS server will only support committing the cached data when the data set is timed out. For the NFS version 4 commit operation, the z/OS Server will only support committing the cached data upon receiving the close operation.

Table 3. Breakdown of text and binary writes

Description	Binary	Text
Data is flushed to DASD when the	Number of cached packets exceeds the amount specified in **cachwindow()**, or the file times out.	File times out. If the number of packets exceeds the amount specified in **cachwindow()**, all new out of sequence packets will be dropped.
Padding	Binary zeros	Binary zeros
Record delimiters	There are no record delimiters. Therefore, there is no attempt to add end of line characters.	There can be record delimiters. Therefore, an end of line character is added to the end of the record.

Case sensitivity–maplower, nomaplower

If the processing attribute **maplower** is specified, the MVS file name is mapped to the lower case when returned to the client. If the processing attribute **nomaplower** is specified in the attributes, all entries in the exports data set are case-sensitive. Therefore the client mount request must specify the MVS qualifier with the correct case to successfully match the exports data set entry.

Selection of text or binary processing modes–text, binary

You can specify either **text** or **binary** processing mode when you access files. This processing mode does not describe the type of data in the original file, but rather, it specifies whether to convert between ASCII and EBCDIC when sending file contents between the z/OS host and the client workstation. See “Mount command syntax and examples” on page 65 for more information about text and binary processing of files using the z/OS NFS client.

Binary processing mode

The **binary** processing mode specifies to send and receive file contents between the z/OS host and the client in binary form, avoiding the ASCII/EBCDIC conversion required in text mode. This is faster than text mode. However, users on MVS cannot read the file, because the contents are not in EBCDIC. Therefore, use the binary processing mode to create or access a file only if the file is not intended to be shared with users on the z/OS host, or the file content is binary.

When fixed-length records are written in binary mode, the server pads the last record of the file with null characters if the last record is less than the fixed record length. These padding bytes are counted in the file size.

Text processing mode

With the **text** processing mode, when data is read, record boundaries are converted to (or from) workstation line terminators such as **lf** or **crlf**, and data representation is changed between EBCDIC and ASCII.

Selection of how blanks are handled–blankstrip, noblankstrip: When fixed-length records are written in **text** mode, records are padded with blanks if the record length is larger than a line, and if the **blankstrip** processing attribute is enabled. (When sending data from MVS, **blankstrip** strips trailing blanks. When sending data to MVS, **blankstrip** pads the records with blanks). In **text** mode processing, data representation is changed from ASCII to EBCDIC, and line terminators are converted to record boundaries. All data is converted according to the active translation table. Therefore, if the data set contains a mixture of characters and binary data, binary data is converted as well. In **text** mode, be careful not to mix your text data (characters) with binary data.

If you are writing data to a fixed-length MVS file in text mode with blank stripping enabled, and the data contains blanks at the end of the line, an I/O error occurs. This is because the server is not able to return the blanks to the client when the file is read back.

If you save a fixed-length MVS file in **text** mode with one or more lines exceeding the maximum record length, an I/O error occurs. For example, suppose an MVS file has fixed-length records of 80 bytes. After you edit the file using the **vi** editor on your workstation, one of the file's records is 83 bytes long (exceeding the fixed length by 3 bytes). When you save the file back to the server, the MVS file is either partially or totally destroyed, and the "I/O Error" message appears on your screen. While you are still in the editing session, save the edited file in an alternate local file. After you correct the local file so that no line exceeds 80 bytes, save it back into the MVS file.

If you get an error message when trying to create or access an MVS file, see Chapter 19, "Network File System messages," on page 241 for further explanation of the message.

Using the lf line terminator: For an AIX or UNIX client, use **lf** as a line terminator when using **text** processing mode.

Number representation

The **text** processing mode does not change the number representation format between the host and client. When you choose **text** as the processing mode, the NFS converts characters between ASCII format and EBCDIC format, and processes line terminators, but no other translation of user data occurs.

When you select **binary** as the processing mode, NFS stores your data unchanged. Therefore, regardless of the processing mode you choose, you cannot change numbers from one client workstation's format to another client workstation's format.

Access to migrated files–retrieve, noretrieve; wait, nowait

Sometimes files on MVS are migrated to another storage level, such as a space-saving format on DASD or tape. If your file has been migrated and you try to access it, it might take a while for it to be recovered back into primary storage. The **retrieve** and **noretrieve** processing attributes control what happens when you try to access a migrated file.

There are three ways that the **retrieve** or **noretrieve** option is controlled.

1. Using the Default Retrieve Attribute

   You can use the default retrieve processing attribute by not entering **retrieve** or **noretrieve** in your **mount** command or **file access** command.

2. Specifying **retrieve** with the **mount** command

   You can issue the **mount** command, specifying **retrieve** or **noretrieve**. The attributes specified in the **mount** command override the attributes in the default attribute data set.

   In this example, migrated files under the mount point are not retrieved. However, you can access files under the mount point which are not migrated.

   ```
   $ mount mvshost1:"smith,noretrieve" /u/smith/mnt
   ```

   Conversely, the next command causes the migrated files under the mount point to be retrieved when accessing the files.

   ```
   $ mount mvshost1:"smith,retrieve" /u/smith/mnt
   ```

3. Specifying **retrieve** with a file access command

   You can issue a file access command with the attribute **retrieve** or **noretrieve** specified. The attributes specification in the **file access** command overrides the attributes in the **mount** command and the server default attributes.

   This command causes all files under the mount point /u/smith/mnt to be retrieved if they are migrated:

   ```
   $ ls -l "/u/smith/mnt,retrieve"
   ```

   This command, however, does not cause migrated files under the mount point /u/smith/mnt to be retrieved:

   ```
   $ ls -l "/u/smith/mnt,noretrieve"
   ```

Access to migrated system-managed data sets

z/OS DFSMS allows access to data set attributes for migrated SMS-managed data sets, without having to recall the data set if the data set was migrated under DFSMS/MVS V1R3 or later. Supported data set types are SMS-managed PS, VSAM ESDS, VSAM KSDS, VSAM RRDS, PDS, and PDSE. Migrated PDS/PDSE members are not supported.

The z/OS NFS server is able to obtain the attributes of a supported SMS-managed migrated data set without recalling the data set. Attributes such as the time stamp and file size are saved to DASD. Subsequent file size requests do not cause a recall of the supported SMS-managed migrated data set, thus improving performance. However, when the data set is modified outside the server by a non-NFS application (for example, by the TSO/E editor) before it was migrated, the stored file size could be incorrect. When the data set is accessed again by the server, a recall must be done to determine the correct file size.

When a request is made to remove any of the supported SMS-managed migrated data sets, the data set will be deleted without recall. For PDS and PDSE migrated data sets, the data set will be recalled in order to read its member information.

File handle refresh

File handles of mounted objects (directories or file systems) are saved on DASD in a mount handle data set and are automatically established again when the server restarts. However, file handles for the files within a mounted object are kept in virtual storage (memory) and they are lost if the server restarts. This may result in

stale file handles for NFS version 2 and 3, or file handle expired in NFS version 4, and the clients may be required to request a new file handle by redoing the **lookup** on the file.

Mapping between the workstation and MVS file systems

In MVS, a file is called a data set, and the two terms are used interchangeably in this book. The z/OS NFS server presents information to you in the form of a UNIX (or AIX) file, though the information is actually stored on MVS-owned DASD in the form of an MVS data set.

The files for a computer system are organized into a file system. The UNIX environment uses a file system which is a hierarchy of directories. MVS, however, uses a non-hierarchical file system in which groups of data sets are referred to by specifying a high-level qualifier.

The MVS high-level qualifier can include the first (leftmost) qualifier of data sets, or the first and second qualifiers, or the first, second, and third qualifiers, and so on. For example, `SMITH` is the high-level qualifier for the files named `SMITH.TEST.DATA` and `SMITH.PROJ7.SCHED`, while `SMITH.TEST` is the high-level qualifier of `SMITH.TEST.DATA` and `SMITH.TEST.DOCS`.

File name extension mapping

File extension mapping allows users to access members of PDS/PDSE data sets on the z/OS host which are mapped from client machine files that contain file extensions. Each PDS/PDSE data set on the host can only be mapped with one unique file extension. For example: IBMUSER.TEXT(M1), IBMUSER.TEXT(M2) will map to m1.txt, m2.txt under directory ibmuser.text on the client machine. This capability allows client machine tools such as editors and compilers, to process host files remotely without modification. There are site and processing attributes that are associated with the file extension mapping. The client user can specify the **fileextmap** attribute to turn the file extension mapping on and the **nofileextmap** attribute to turn the file extension mapping off.

The rules for file extension mapping are contained in a data set called a *side file*. You can establish a default side file, with the default rules for file extension mapping, by specifying the **sidefile(dsname)** attribute in the attributes data set. A client user can also specify this attribute on a mount command to override the default side file name, as shown in the following example:

```
[C:\] mount z: mvshost1:"user1.pds,sidefile(hlq.nfs.mapping)"
```

The side file specified at the **mount** command will be searched first followed by the default side file. The system administrator can specify the maximum space available for side files using the site attribute **sfmax(n)**. A sample mapping side file is provided as GFSAPMAP in the NFSSAMP library. See “Processing attributes syntax” on page 93 for more information.

Specifying the side file on an NFS version 4 **mount** command has the following effects:

- Loads the side file if it is not loaded
- Does not reload the side file if it was already loaded (does not change the current side file).

Note: With NFS version 4, an **unmount** command does not unload the side file, because no UNMOUNT_RPC is sent to the z/OS NFS server.

Mounting of MVS data sets onto a client mount point

To access an MVS file system from the client, you use the **mount** command to create a temporary link between specific MVS data sets and a UNIX directory (preferably empty) or an unused logical drive. The UNIX directory or drive is called a mount point.

You use an MVS high-level qualifier in the mount command to specify which MVS data sets to mount onto a mount point. The MVS data sets beginning with the specified high-level qualifier appear as files under the mount point. See Figure 2 on page 33.

You can also perform a mount using a fully qualified data set name or an alias to a user catalog, but not the catalog name itself. Only cataloged data sets are supported by the z/OS NFS server, and tape data sets are not supported.

Data set organizations supported include:

- Physical sequential (PS)
- Direct access (DA)
- Partitioned data sets (PDS)
- Partitioned data sets extended (PDSE)
- VSAM KSDS
- VSAM ESDS
- VSAM RRDS
- SAM extended format data sets

Both SMS-managed and non-SMS-managed data sets are supported.

- NFS supports multivolume data sets.
- Generation data sets are not supported.

Notes:

1. The filesize for the MVS conventional data set as a directory has a dummy size with a value of 8192.
2. For the NFS Version 3 CREATE procedure, the z/OS server does not harden the exclusive create verify token to disk.

Mount examples

Table 4 shows how to mount z/OS MVS data sets from various platforms.

Table 4. Examples of mounting MVS data sets from clients

Clients	Command Examples
UNIX	mount -o vers=n,sec=r mvshost1:MVSHLQ,"procattr1,procattr2,...." /u/smith/mnt
Linux (NFS version 4)	mount -t nfs4 -o sec=r mvshost1:MVSHLQ /u/smith/mnt
Linux (NFS version 2 and 3)	mount -o vers=*n* mvshost1:MVSHLQ /u/smith/mnt
Windows	nfs link z:'\\mvshost1\MVSHLQ,procattr1,procattr2,...' id pw

In the examples:

Operand	Description
mvshost1	Specifies the name of the z/OS host.

MVSHLQ	Specifies the high-level qualifier of the MVS data set.
procattr	Specifies any valid processing attribute, such as text/binary or hfsfbtimeout(n).
/u/smith/mnt	Specifies the local mount point.
z:	Specifies the drive letter on the Windows system.
id pw	Specifies login id and password for pcnfsd.
-t nfs4	Specifies NFS protocol version 4 for Linux (optional).
-o vers=n	Specifies the NFS protocol version to be used (2 or 3 for Linux; 2, 3 or 4 for others) (optional).
-o sec=r	Specifies RPCSEC_GSS security flavors, which is available only on the z/OS NFS version 4 server. Valid options are **sys**, **krb5**, **krb5i**, and **krb5p**.

Variants of the mount command

The name of the command that performs the mount operation varies for different client implementations. For example, the Windows implementation uses the **nfs link** command while most UNIX and Linux implementations use the **mount** command.

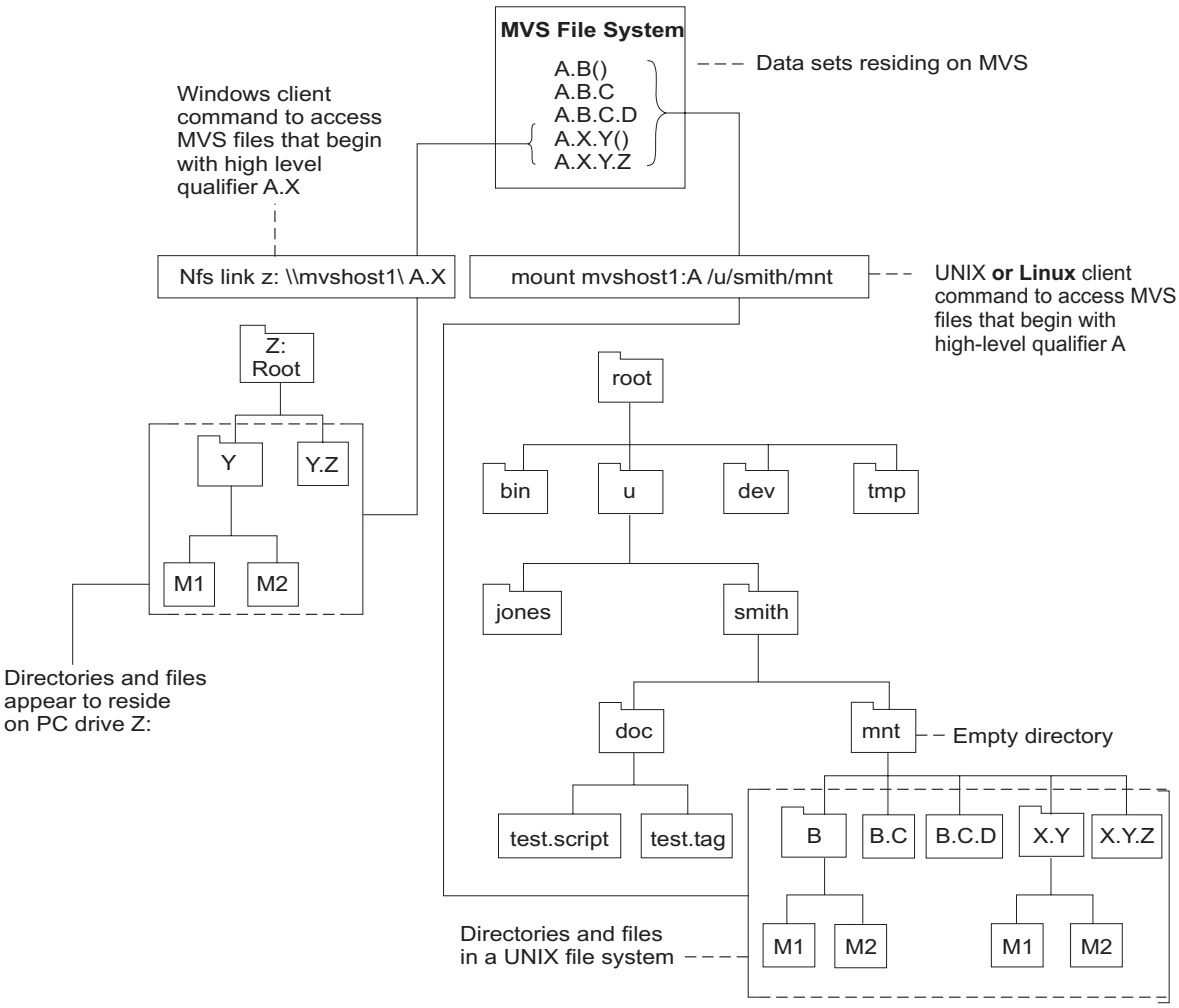

Figure 2. Examples of mounting MVS data sets on Windows, UNIX and Linux clients. The Windows client is mounting the MVS data sets which begin with the high-level qualifier "A.X". The UNIX and Linux clients are mounting the MVS data sets which begin with the more inclusive high-level qualifier "A". The parentheses indicate a PDSE containing the members M1 and M2. The PDSEs A.B() and A.X.Y() appear as directories, and their members appear as files within those directories.

Use of a PDS or PDSE as a directory

If the data sets specified include partitioned data sets, a second level of hierarchy is shown. This allows you to define one level of directories under the mount point. Thus, you can issue **mkdir** to create a directory (stored as a PDS or PDSE) and then create files (stored as members of a PDS or PDSE) within that directory.

This use of a PDSE is shown in Figure 2, which illustrates the mapping of file names between client file systems and the MVS file system resulting from a **mount** command.

Use of multiple mount points

You can arrange groups of data sets into several UNIX mount directories (or PC mount drives) by using MVS naming conventions to mount specific data sets at each mount point. For example, you could mount user1.project1 to get all data sets beginning with user1.project1 mounted at one point in the local file system, and you could mount user1.project2 at another point. This would create the effect of

two distinct directories (or drives), one containing the user1.project1.* data sets, and the other containing the user1.project2.* data sets.

Data set serialization and sharing

The z/OS NFS server handles data set serialization and sharing differently depending on the type of data set:

Physical sequential
: The server insures the read/write integrity of a physical sequential data set by SVC 99 dynamic allocation with exclusive option whenever a physical sequential data set is opened for output; otherwise it is allocated with share option.

VSAM data set
: The server dynamically allocates a VSAM data set with share option and allows the VSAM access method to manage data sharing via the SHAREOPTIONS specified during data set definition.

PDSE data set
: The server dynamically allocates a PDSE data set with share options and allows the PDSE functions to manage the integrity of the PDSE data set and its members.

PDS data set
: The server dynamically allocates a PDS data set with share option and surrounds the PDS and its members with exclusive ENQs against the QNAME=SPFEDIT and RNAME=data set name. This does not protect the PDS from other z/OS users who are attempting to access the PDS without performing ENQ against SPFEDIT similar to the z/OS NFS server.

NFS protocol

Table 5 illustrates that the NFS procedures are not all supported for conventional MVS data sets.

Table 5. NFS procedures

Procedure	Version 2 Protocol	Version 3 Protocol	Version 4 Protocol
link	no	no	no
mknod	N/A	no	no
readlink	no	no	no
readdirplus	N/A	no	no
setattr	yes	yes	yes
statfs	yes	N/A	N/A
symlink	no	no	no

Note: Setattr only supports filesize=0 truncation and UNIX permission is set to 777.

NFS file system attributes

The z/OS NFS server generates MVS-specific values for certain UNIX file system attributes. Table 72 on page 321, Table 73 on page 321, and Table 74 on page 321 illustrate the MVS values that the z/OS NFS server generates.

Chapter 4. Using z/OS UNIX System Services files

This chapter explains what you need to know to access z/OS UNIX files from a client workstation.

- The z/OS UNIX file system
- POSIX compatibility
- Attributes specific to z/OS UNIX
- Protecting your z/OS UNIX files
- Accessing z/OS UNIX files from the client
- Linking an MVS data set to a hierarchical file system
- UNIX look and feel

Note: For detailed information about z/OS UNIX, see *z/OS UNIX System Services User's Guide*.

z/OS UNIX file system

z/OS UNIX provides a hierarchical file system (HFS) for z/OS. z/OS UNIX also provides the z/Series File System (zFS). An HFS or zFS file within z/OS UNIX is called a z/OS UNIX file. HFS and zFS files are organized in a hierarchy of files and directories in a tree structure. A directory can contain files or other sub-directories. The highest level directory is called the root directory. Figure 3 shows an example of mounting an HFS or zFS directory from a UNIX client.

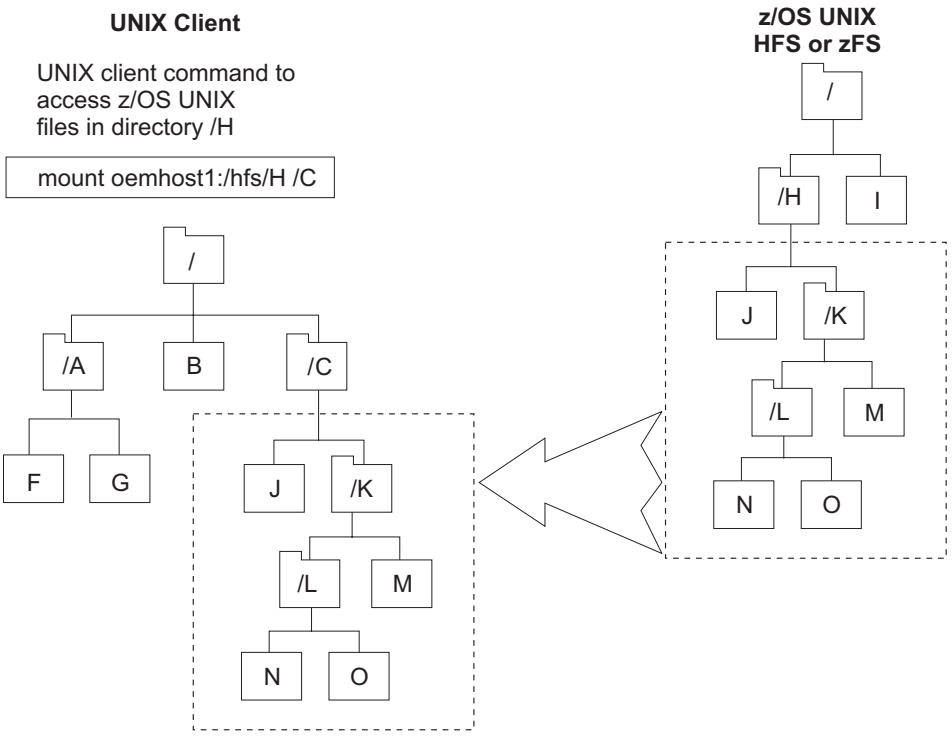

Figure 3. Example of mounting an HFS or zFS file from a UNIX client

© Copyright IBM Corp. 1991, 2006

A z/OS UNIX file system must be mounted by an MVS system operator using a TSO MOUNT command before that z/OS UNIX file system can be mounted by an NFS client via the z/OS NFS server.

z/OS UNIX files are byte-oriented rather than record-oriented (unlike conventional MVS data sets). This data can be shared with TSO/E z/OS UNIX users in addition to NFS clients. All data written to a z/OS UNIX file can be read by all programs as soon as it is written. You can also copy data between z/OS UNIX files and MVS data sets using z/OS UNIX utilities like ISHELL.

POSIX compatibility

The NFS supports file access to the z/OS UNIX file system. z/OS UNIX supports a set of standards called the portable operating system interface (POSIX). See *z/OS UNIX System Services User's Guide* for more information about POSIX compliance. With z/OS UNIX, the NFS performs the following functions.

- Supports hierarchical directories
- Allows file names up to 255 characters in length
- Allows path names up to 1023 characters in length
- Supports mixed-case names and special characters, except nulls, slashes, and commas in file and path names
- Supports UNIX-style file access permissions
- Supports group ID and user ID at the file level
- Supports the full NFS protocol (including external links)
- Enables data sharing between clients and the z/OS UNIX
- Enables you to link conventional MVS data sets to a POSIX path name

This support incorporates the basic strengths of the z/OS system for both existing MVS data and applications and for new POSIX conforming data and applications.

NFS protocol

z/OS UNIX is compliant with all of the z/OS Network File Systems version 2, version 3 and version 4 protocols.

Attributes specific to z/OS UNIX System Services

The following attributes are specific to the z/OS UNIX:

sync	Processing attributes for version 2 protocol only
async	Processing attributes for version 2 protocol only
hfs	Site attribute
extlink	Attribute, see "Linking an MVS data set to a z/OS UNIX file system" on page 40

Note: These attributes are also explained in Chapter 9, "Initialization attributes for the z/OS NFS server," on page 87.

Synchronous write to a z/OS UNIX file for NFS version 2 protocol

Use the **sync** and **async** processing attributes to specify whether data received by a write request for a z/OS UNIX file object is committed to nonvolatile storage before the write response is returned to you.

If **sync** is specified for a z/OS UNIX file object, the data is written to z/OS UNIX and immediately committed to non-volatile storage.

For greater throughput, you can alternatively specify **async**. Your data is then committed to the disk some time after the write request is received from the NFS client. Your data is written to disk when the write timeout occurs, or if z/OS UNIX reclaims buffer cache storage.

The **sync** and **async** processing attributes only apply to z/OS UNIX data access. They are ignored for MVS data set access. A TSO/E z/OS UNIX user doesn't have to wait for the data to be committed to nonvolatile storage before accessing. z/OS UNIX maintains a central buffer cache and a TSO/E z/OS UNIX user can access the data as if it were in the file once it has been written by the z/OS NFS server.

Synchronous write to a z/OS UNIX file for NFS version 3 or 4 protocol

For the NFS version 3 WRITE procedure or NFS version 4 WRITE operation there is a processing argument ***stable*** and output parameter ***commit*** which specifies whether data received by a write request for a z/OS UNIX file object is committed to nonvolatile storage before the write response is returned to you.

If the ***stable*** processing argument is used during the write procedure, there are three modes when the write procedure writes to a file:

file_sync
: The z/OS NFS server must commit all data written and all file system data to stable storage before returning commit results.

data_sync
: The z/OS NFS server must commit all data written and sufficient metadata to enable retrieval of data, before it returns a reply to the client.

unstable
: The z/OS NFS server may not commit any part of the data and metadata to stable storage, before returning a reply to the client. The data will be committed when a timeout occurs.

For the commit procedure, the z/OS server will support committing the entire data and metadata to stable storage.

Authorization checking when writing to a z/OS UNIX file

The z/OS Network File System server allows the owner of an z/OS UNIX file to write to the file regardless of the UNIX permission bits setting on the file.

HFS site attribute

To access z/OS UNIX files, you must know the HFS prefix defined by your system administrator (the default is /hfs). You can use the **showattr** command to display the HFS prefix defined for your location. You use this prefix in your mount command before the path name of the z/OS UNIX directories that you are

mounting. The HFS prefix is used by the NFS server to distinguish z/OS UNIX directories from conventional MVS data sets, but the HFS prefix isn't part of the path name that you see. After you have entered the mount command, you access HFS files using the local mount point.

Protecting your z/OS UNIX System Services files

As an z/OS UNIX user, you can control the read, write, and execute access to your files by other users in and outside of your group by setting the permission bits associated with the files.

To access z/OS UNIX files from the NFS, you must be defined as an z/OS UNIX user. The system programmer defines you as an z/OS UNIX user by assigning an z/OS UNIX user ID (UID) and an z/OS UNIX group ID (GID) to you. The UID and GID are numeric values associated with a TSO/E user ID. The values are set in the RACF user profile and group profile when you are authorized to use z/OS UNIX. The system uses the UID and GID to identify the files that you can access. Your specific UID value identifies you as a user of z/OS UNIX services. A GID value is a unique number assigned to a group of related users. These numbers appear in the RACF user profile. See *z/OS UNIX System Services Planning* for more information.

Accessing z/OS UNIX System Services files from a client

Most of the commands that are used to access z/OS UNIX files are identical to the following commands that are used to access conventional MVS data sets.

mvslogin
showattr
mount
umount
mvslogout

The only command that is changed for z/OS UNIX is the **mount** command.

Note: The syntax of the commands listed above may vary between platforms, see the appropriate chapter for examples specific to the platform you are using to access z/OS UNIX files.

If you are using AIX (or any other UNIX-based operating system) see Chapter 6, "Commands and examples for AIX and UNIX clients," on page 47.

Mount examples

Table 6 shows how to mount z/OS UNIX files from various platforms.

Table 6. Examples of the mount command for clients

Clients	Command Examples
AIX, UNIX, Solaris	mount -o vers=*n*,sec=r mvshost1:"/hfs/smith" /u/smith/mnt
Linux with NFS V4	mount -t nfs4 -o sec=r mvshost1:"/hfs/smith" /u/smith/mnt
Linux with NFS V2 or V3	mount -o vers=*n* mvshost1:"/hfs/smith" /u/smith/mnt
Windows	nfs link z: '\\mvshost1\hfs/smith' /n id pw

In the examples:

Operand	Description
mvshost1	Specifies the name of the MVS host.
/hfs	Specifies the HFS prefix.
/smith	Specifies the HFS directory to be mounted.
/u/smith/mnt	Specifies the local mount point.
-t nfs4	Specifies NFS protocol version 4 for Linux (optional)
-o vers=n	Specifies the NFS protocol version to be used (2 or 3 for Linux; 2, 3, or 4 for others) (optional)
-o sec=r	Specifies RPCSEC_GSS security flavors, which are available only on the z/OS NFS version 4 server. Valid options are**sys**, **krb5**, **krb5i**, and **krb5p**.
/n	Specifies the NFS protocol version for Windows (2, 3, or 4) (optional)
id pw	Specifies login id and password for pcnfsd.

Note: The */hfs* prefix value is used by the z/OS NFS server to determine if a file is a z/OS UNIX file, and does not appear in the path name of an HFS file once it is mounted.

z/OS UNIX data transfer and conversion

The transfer and conversion of z/OS UNIX data between the client and server is handled differently depending on the version of the NFS protocol in use.

Data transfer under the NFS version 4 protocol

With the NFS version 4 protocol, text data and metadata are transferred between the server and client in the **UTF-8** data format. With NFS version 4, ASCII text is not transferred directly.

Text or binary processing - NFS version 2 and 3 protocols

z/OS UNIX is a byte-oriented, hierarchical, EBCDIC file system. For the NFS version 2 and version 3 protocols, the z/OS NFS server provides ASCII to EBCDIC text translation. For these versions of NFS, if you are just using the mainframe as a repository for your workstation (ASCII) data, you should use the binary mode to speed processing. If you use text mode, data from your workstation is converted into EBCDIC when it is stored on the mainframe. Conversely, when the z/OS NFS server returns the data to your client system, it converts the data back into ASCII. The conversion can slow processing, but might be necessary if you are sharing data with other MVS users. All data is converted according to the active translation table. Therefore, if the data set contains a mixture of characters and binary data, binary data is converted as well. In text mode, then, be careful not to mix your text data (characters) with binary data.

In text mode, you can either use the OEMVS311 translation table or the replacement customized translation table to convert data between ASCII and EBCDIC. If you are using z/OS UNIX and text mode processing, specify the OEMVS311 translation table with the xlat processing attribute. The OEMVS311 table translates ASCII (ISO 8859-1) to and from EBCDIC (1047 - z/OS UNIX System Services). TCP/IP for MVS version 3.1 provides the OEMVS311 table. This

table translates the UNIX line terminator (lf) to the z/OS UNIX line terminator (nl). See *z/OS Communications Server: IP Configuration Reference*for more information about creating and customizing your own translation tables.

This is an example of specifying the OEMVS311 translation table during a mount:

```
$  mount 1stc3mvs:"/hfs/usr/man/C,text,xlat(oemvs311)" /mnt
$  export MANPATH=/mnt
$  man more
```

In this example:

Operand	Description
1stc3mvs	Specifies the name of the z/OS host.
/hfs	Specifies the HFS prefix.
/urs/man/C	Specifies the HFS directory to be mounted.
text	Specifies that data be converted between ASCII and EBCDIC
xlat(oemvs311)	Specifies that the translation table named OEMVS311 is used to convert data between ASCII and EBCDIC.
/mnt	Specifies the local mount point.
man more	Obtains a Man Page description of the **more** command

Linking an MVS data set to a z/OS UNIX file system

This section explains how to access an MVS data set via a z/OS UNIX path name by using the external link command. It also explains how to display the contents of an external link and how to delete an external link.

Creating an external link

You can create an external link to an MVS data set, and then transparently access the MVS data set by referencing the external link. The external link simulates a UNIX-like hierarchical naming convention for conventional MVS data sets. This is done using the **ln** command, for example:

```
mount mvshost1:USER1 /mnt
mount mvshost1:/hfs/u/nfs /samples
ln -s USER1.MVSFILE /samples/linkfile,"extlink"
```

In this example a z/OS UNIX file object, */linkfile*, of the file type "extlink" is created containing the file name of the MVS data set `USER1.MVSFILE` to be accessed. The source file must be mounted to a z/OS UNIX file system. The external link must reference an MVS data set. All future references to */samples/linkfile* access `USER1.MVSFILE` transparently.

In this example the file */usr/pub/myfile* is copied to the MVS data set `USER1.MVSFILE` that is contained in the external link */samples/linkfile*:

```
cp /usr/pub/myfile /samples/linkfile
```

Your installation should make sure that the appropriate security permissions have been obtained to access the MVS data set. You will receive "Permission Denied" message if the mount point */mnt* has not been established on USER1.

A mount point must be established before the external link is established. Otherwise, the error code ACCESS DENIED is returned. For physical sequential data sets, the high level qualifier of a data set must be established. For example, if

you had a file called `smith.test.data` you can mount with `smith`, `smith.test`, or `smith.test.data` as your high level qualifier. For PDS and PDSE data sets, the fully qualified name must be established as a mount point. An example of a fully qualified name would be, `smith.test.data`.

Displaying the contents on an external link

You can display the contents of an external link by appending the ″extlink″ sequence to the external link path name. This permits the user to inspect the contents of the external link with the **ls -l** command.

This example shows how to display the attributes and contents of the external link `/samples/linkfile`:

```
ls -l /samples/linkfile,"extlink"
lrwxrwxrwx  1 user1 13 Jun 17 20:43 /samples/linkfile ->USER1.MVSFILE
```

This example shows how to display the attributes of the MVS target data set `USER1.MVSFILE`:

```
ls -l /samples/linkfile
-rw-rw-rw-  1 root  2112 Sep 28 13:50 /samples/linkfile
```

Deleting an external link

The external link file object is deleted with the remove request:

```
rm /samples/linkfile,"extlink"
rm /samples/linkfile
```

Either **rm** command results in the z/OS UNIX external link file object alone being removed. The target MVS data set is not affected.

UNIX look and feel

Using the z/OS NFS with z/OS UNIX managed files provides UNIX client users with a transparent view of their data. The file attributes are maintained in the same way as is found on any UNIX system.

- Regular, directory, link, device, and FIFO file types
- User, group, and other read/write/execute access permissions
- UID and GID file ownership
- File size

To access z/OS UNIX files, it is necessary to be defined to RACF as an z/OS UNIX user. Some installations might prefer to provide users with unrestricted access to their z/OS UNIX data by specifying `security(none)` or `security(exports)` in the site attributes. With this setting, the client's user ID and group ID credentials are used for all file access authentication, and there is no requirement for the user to be defined to RACF or to perform the mvslogin command.

Note: For `security(none)` and `security(exports)` options, the UID of the root (UID=0 from the workstation) is mapped to UID of NOBODY (UID=65534) by the z/OS NFS server. The implication is that the z/OS NFS server will use the mapped UID of 65534 for all z/OS UNIX file authorization checking. For example, file creation owner UID is set to 65534 in the z/OS UNIX file attribute.

NFS file system attributes

For z/OS UNIX files, see Table 75 on page 322 and Table 76 on page 322 for the file system values that are returned for NFS attributes.

Chapter 5. Locking and access control

This chapter provides an overview of the z/OS NFS locking and access control functions provided by the Network Lock Manager (z/OS NFS NLM) and the z/OS NFS Network Status Monitor (z/OS NFS NSM). It explains how they work together to provide file locking and access control capability over z/OS NFS. In addition, this chapter also explains the following features:

- Monitored lock
- Non-monitored locks
- Locking files
- Locking records

In z/OS V1R7, NLM and NSM are integrated into the z/OS NFS server to facilitate the expanded locking and serialization functions. Separate procedures for starting and stopping NLM and NSM are replaced by the server site attribute **nlm | nonlm**, which specifies their startup along with the NFS server. This integration also coordinates the locking function with stale file handle processing; when a file handle becomes stale, not only will the code clean up the file related blocks as it does in prior releases, but it will also release any locks which remain held for that file.

Using Network Lock Manager (NLM) in NFS V2 and V3

In NFS version 2 and version 3, the z/OS NFS NLM allows a client on the host to lock a record or a file on the z/OS NFS server. A client user can either choose to lock the entire file or a record section of a file. The two types of locks that the client host uses are monitored locks and non-monitored locks.

The z/OS NFS NLM only supports advisory locking. Advisory locking is when the operating system keeps track of which files have been locked by which process, but does not prevent a process from writing to a file that is locked by another process. This means that a process can ignore an advisory lock if the process has adequate permission.

Monitored locks

Monitored locks provide the client user with reliability. If the server host on which the monitored locks are established fails, the locks are reinstated when the server host recovers. The locks that are held by the client host are discarded by the z/OS NFS NLM on the server host if the client host fails before the locks are released. Monitored locks will only work correctly if both the server host and the client host are running NSM.

Non-monitored locks

Non-monitored locks are used on personal computer (PC) operating systems. Non-Monitored locks provide the same functionality as the monitored locks with one exception. If the server host on which the locks are established, fails and recovers, the locks will not be re-established. The client host is responsible for detecting a server host failure and re-establishing the locks. In addition, the client host informs the z/OS NFS NLM when it has rebooted so that the client host can discard all of the locks and file shares held for the client.

© Copyright IBM Corp. 1991, 2006

Specifying a grace period for reclaiming locks

You can specify a time limit, or grace period, for clients to reclaim NFS V4 or NLM locks and share reservations when the z/OS NFS server restarts after a failure. To set this time limit, use the **leasetime** site attribute. For details, see "Site attributes syntax" on page 104.

During the reclaim grace period after a z/OS NFS server restart, the grace period may be extended after an open or lock reclaim event. If the **leasetime** site attribute is greater than 1200 seconds (20 minutes), the grace period will not be extended at a reclaim event; if the **leasetime** site attribute is less than 20 minutes, the grace period will be extended to one lease time after the reclaim event, up to but not exceeding 20 minutes after the z/OS NFS server completed its restart.

Listing locks held for a file

To diagnose possible problems with conflicting locks, z/OS operators can issue a **listlock** command that displays all client programs and users which hold a lock on a file. The output messages include client and user id, the lock ranges held, and lock status. The **listlock** command can be used for MVS data sets, PDS or PDSE members, or z/OS UNIX files. For more information, see "Listlock operand" on page 170.

You can use the listlock command to find locking information in cases where a lock is unavailable and the blocker is managed by another NFS server address space running on the z/OS system. To determine the identity of the blocker in this case, the listlock command should be issued on the system which owns the lock.

You can also detect some situations in which a deadlock occurs. Use the z/OS UNIX System Services **v_lockctl** service to detect deadlocks on both MVS data sets and z/OS UNIX files on the z/OS NFS server, when the deadlocks are under the scope of z/OS UNIX locking services.

Note: The NLM protocol does not provide any means for the server to notify the client that its lock has been released. Thus use of this operator command can lead to data integrity problems.

Releasing locks held for a file

To release all locks for a file, z/OS operators can issue a **release** command that releases locks for z/OS UNIX files and MVS data sets or members. The command also forces the NFS server to release the file, and if the file is active, to close and deallocate it. For more information, see "Release operand" on page 167.

Note: The NLM protocol does not provide any means for NFS server to notify the NFS client that the locks were released. Thus, the client may proceed under the false assumption that it still has the locks. Therefore, the release command should only be used in extreme circumstances.

Using Network Status Monitor (NSM) in NFS V2 and V3

In NFS version 2 and version 3, the z/OS NFS NSM is a service that provides applications with information on the status of network host. Each z/OS NFS NSM keeps track of its own ″state″ and notifies any interested parties of a change in its state.

For correct operation of the z/OS NFS NSM, the client and the server hosts are required to monitor each other. When a lock request is issued by a process running on the client host, the NLM on the client host requests the NSM on the client host to monitor the server host. The client NLM then transmits the lock request to the z/OS NFS NLM on the server. On reception of the lock request the z/OS NFS NLM on the server host will request the z/OS NFS NSM on the server host to monitor the client host. In this way, each host is monitored by the NSM on the other host.

Chapter 6. Commands and examples for AIX and UNIX clients

This chapter gives the syntax and examples of commands that AIX users need to know to access MVS data sets from a client. Some examples are also provided for UNIX and Sun Solaris environments. This chapter shows how to perform the following tasks.

- Log on to z/OS from your client
- Access MVS data sets from your client
- Display default mount point attributes
- Query mount points
- Unmount MVS data sets from the client
- Log out of z/OS.

The **mount** and **umount** commands are operating system specific commands. They are not shipped with z/OS NFS. See your NFS client documentation for the exact syntax and usage.

Using commands on AIX

Figure 4 shows a summary of the syntax for the commands described in this chapter:

The **mvslogin** command is used to log in to z/OS from your workstation. The mvslogin command can be issued multiple times, and the last one overrides the previous one. The mvslogin command is only required when accessing data on systems where the z/OS NFS server site ***security*** attribute is set to *saf* or *safexp*.

The following is the mvslogin command syntax.

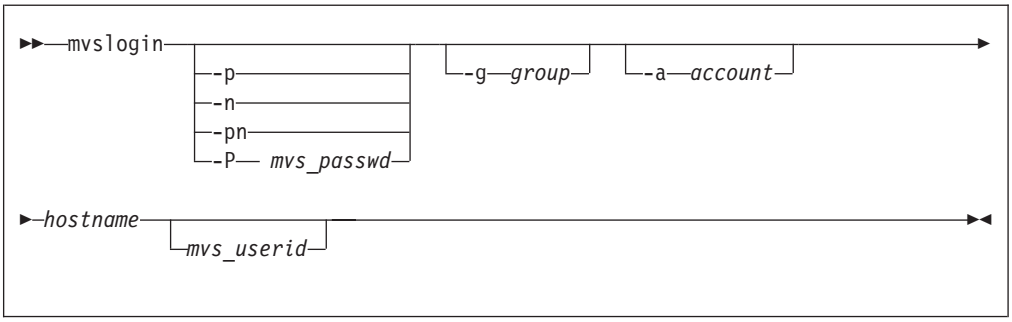

Figure 4. Syntax for commands

where

Operand	Description
-p	Causes a prompt for your z/OS password. The password is passed to z/OS to validate the user logging in. Your security procedures determine whether you should use this parameter.
-n	Causes a prompt for a new password.
-pn	Causes a prompt for the user's current password and then causes two prompts for the user's new password.

© Copyright IBM Corp. 1991, 2006

-P *mvs_passwd* No prompt for your z/OS password; just type your z/OS password after the -P. This enables you to automate your z/OS login.

-g *group* A group name string passed to z/OS for accounting purposes. The maximum length is 8 characters.

-a *account* An account string passed to z/OS for accounting purposes. The maximum length is 16 characters.

hostname The name of the z/OS host (for example, `mvshost1`).

mvs_userid A user ID that z/OS recognizes as valid. If you do not specify this parameter, your workstation user name is used. The z/OS NFS server does not support the use of an alias user ID or a mixed case user ID with the mvslogin command.

The **mount** command is used to make a connection between a mount point on your local file system and one or more files in the z/OS file system.

The following is the **mount** command syntax.

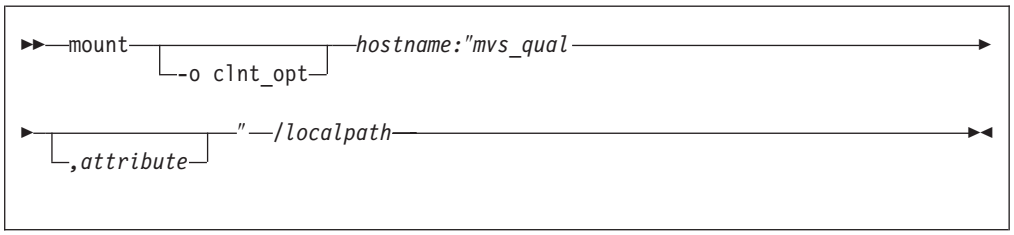

where

Operand	Description
-o *clnt_opt*	The client mount command options (such as `soft,timeo=20`). Refer to the documentation of your client operating system for a description of the options for your client environment.
hostname	The name of the z/OS host (for example, `mvshost1`).
mvs_qual	The path name of an z/OS UNIX file (must begin with the /HFS prefix) or the high-level qualifier(s) of a conventional MVS data set that you are accessing in the z/OS file system (such as `smith` or `smith.project`).
attribute	A z/OS NFS server data set creation or file processing attribute (such as `text`). See Chapter 9, "Initialization attributes for the z/OS NFS server," on page 87. If you specify any attributes, make sure you enclose *mvs_qual* and the attributes in double quotation marks.
/localpath	The mount point on your client system (for example, `/u/smith/mnt`). This should be an empty directory.

The **showattr** command is used to display the default attributes or the attributes that have been set for a specific mount point. If you specify a mount point, showattr shows the attributes for the mount point, including the overriding values. For descriptions of the attributes, see Chapter 9, "Initialization attributes for the z/OS NFS server," on page 87.

The following is the **showattr** command syntax.

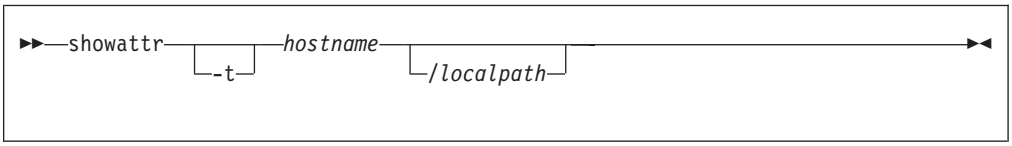

where

Operand	Description
-t	Used to specify tersed output.
hostname	The name of the z/OS host (for example, `mvshost1`).
/localpath	The mount point on your client system (for example, `/u/smith/mnt`). This should be an empty directory.

The **umount** command is used to break the connection between the mount point on your client and the server. When you issue this client command, the file you were editing is released (written to DASD). You do not need to unmount after each session, unmount only when you no longer have a need to access the z/OS file system. Check the documentation for your client operating system to ensure that you enter the umount command correctly.

The following is the **umount** command syntax.

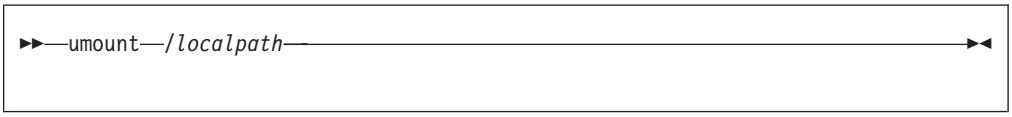

where

Operand	Description
/localpath	The mount point on your client system (for example, `/u/smith/mnt`). This should be an empty directory.

The **mvslogout** command is used to disconnect from the remote z/OS NFS server host. The mvslogout command is only required when the mvslogin command was used to begin the connection.

The following is the mvslogout command syntax.

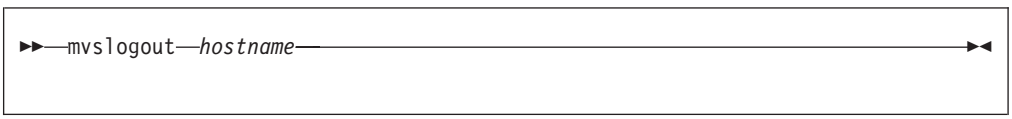

where

Operand	Description
hostname	The name of the z/OS host (for example, `mvshost1`).

Quick reference of AIX and UNIX commands

The following information is an example of a standard z/OS login and logout procedure for AIX.

```
# mvslogin mvshost1 smith
Password required
GFSA973I Enter MVS password: password
GFSA955I smith logged in ok.
# mount mvshost1:"smith" /u/smith/mnt
mount: mvshost1:"smith"
"smith" was attached successfully.
# umount /u/smith/mnt
Unmounting '/u/smith/mnt' ...    successful
# mvslogout mvshost1
UID 200 logged out ok.
```

In this example:

Operand	Description
smith	Specifies a z/OS user ID and high level qualifier for MVS data sets.
mvshost1	Specifies the system name of the z/OS host.
/u/smith/mnt	Specifies the name of the mount point.
GFSA*nnnt*	Messages starting with GFSA apply towards z/OS NFS requests. These messages are explained in Chapter 19, "Network File System messages," on page 241.

You can use the **mount** command with no parameters specified to list the mount points on your client system.

Accessing z/OS UNIX System Services and conventional MVS files

To access z/OS UNIX files or conventional MVS data sets, enter both the **mvslogin** command to log in to the z/OS host system and the **mount** command to mount the files or data sets to your local system. The mvslogin command is only required when accessing data on systems where the z/OS NFS server site security attribute is set to *saf* or *safexp*. Once the files or data sets are mounted to a local directory, you can read, write, create, delete, and treat the mounted files as part of your workstation's local file system. When you are finished with your work, use the **umount** and **mvslogout** commands to break the connection. The mvslogout command is only required when the mvslogin command was used to begin the connection.

To access files on z/OS systems where the z/OS NFS server site security attribute is set to *saf*, *exports*, or *safexp*, you need a z/OS user ID and password, and authorization to access the files that you need. You can only establish one z/OS session for each z/OS user ID. If you do not already have a z/OS user ID, a z/OS password, and access authorization, request them from your z/OS system administrator.

Note: If you cannot use the mvslogin, mvslogout, or **showattr** commands, they might be installed incorrectly or in another directory. Ensure that your system administrator has made the executable code for these three commands available to your workstation and that you have been given the correct path name to find the commands. Also, make sure that your version of these commands matches the release of the z/OS NFS that you are using. Otherwise, the commands might not function properly.

Mvslogin command examples

Use the **mvslogin** command to log in to z/OS from your workstation. The mvslogin command can be issued multiple times and the last one overrides the previous ones. The mvslogin command is only required when accessing data on systems where the z/OS NFS server site ***security*** attribute is set to *saf* or *safexp*.

Table 7 shows examples of the mvslogin command where `mvshost1` is the name of the z/OS host and `smith` is the user's ID on z/OS.

Table 7. Examples of the mvslogin command for clients

mvslogin -p mvshost1 smith
mvslogin -P smithspw -g finance -a 5278 mvshost1 smith
mvslogin -n mvshost1 smith
mvslogin -pn -a 5278 mvshost1 smith
mvslogin mvshost1
mvslogin mvshost1 smith

In the example where the user enters mvslogin mvshost1, the current login client user ID is used as the z/OS user ID.

In the example where the user enters mvslogin mvshost1 smith, the system then prompts for Smith's z/OS password. If Smith logs in successfully, this message appears:

```
GFSA955I smith logged in ok.
```

Otherwise, an appropriate error message appears.

Note: Messages with the prefix of GFSA and GFSC apply to NFS requests. These messages are further explained in Chapter 19, "Network File System messages," on page 241.

When an z/OS UNIX UID or GID segment is defined with the user identification, an additional check is done to compare the z/OS UNIX UID or GID with the client UID or GID during the login processing. An informational message is returned when the server and the client UID or GID do not match. This informational message contains the z/OS UNIX UID and GID for the z/OS user identification.

Note: The authentication is considered successful even if the UID and GID do not match. The message is issued for the user's information only.

For the PCNFSD authentication request, the server UID and GID is returned to the client user if the UID and GID are defined. Otherwise, an arbitrary number is generated and returned to the client user.

"Permission denied" message

If you have successfully logged in and get the "Permission denied" message while trying to access data, that can be due to one of the following cases:

- An **mvslogout** command for the same z/OS host has been entered from the same client platform. See mvslogout for details.
- Your z/OS user ID has been automatically logged out because the **logout** attribute value has been exceeded. This can happen when you leave the client idle for too long. Enter the **mvslogin** command again, and start your processes

again. To find out how many seconds you can stay logged in while your client is idle, issue the **showattr** command and look at the logout attribute.
- Another mvslogin to the same z/OS host using the same z/OS ID has been entered from the same UID and the same client platform. If this is the case, retry your access.

For more information, including additional causes of this message, see the explanation of the ″Permission denied″ message in “Messages from the client platform (AIX)” on page 295.

Note: Some clients give a somewhat different message such as “Access is Denied”.

Mount command examples

Use the **mount** command to make a connection between a mount point on your local file system and one or more files in the z/OS file system.

```
# mount mvshost1:"smith" /u/smith/mnt
```

In this example:

Operand	Description
mvshost1	Specifies the name of the host server.
smith	Specifies the name of the high-level qualifier of the MVS data sets.
/u/smith/mnt	Specifies the name of the mount point (preferably an empty directory).

At the same time, you can also predefine some of the attributes for the mount point, overriding the default attributes.

Overriding default attributes

To override the default attributes, specify different attributes with the **mount** command or in a file access or creation command (such as **vi** in AIX or UNIX).

There are two kinds of attributes that you can modify:

Data set creation attributes provide information about an MVS data set to the z/OS NFS server, such as:
- The type of data set
- How the data set is allocated

Note: Data set creation attributes do not apply to z/OS UNIX files.

Processing attributes provide information to the z/OS NFS server about how to handle the file. For example:
- How long the files remain open
- Whether the file contents are sent and received in text form, or in binary form to avoid ASCII/EBCDIC conversion

Use the **showattr** command to display the default data set creation and processing attributes. For descriptions of the attributes, see Chapter 9, “Initialization attributes for the z/OS NFS server,” on page 87.

Files are created and processed using the mount point attributes that were in effect when the files were last mounted. If your installation's default attributes have been changed (by way of the *exportfs* operand of the **modify** system operator command or restart of the server) and you want to apply these new default attributes, you can unmount and remount (using the **umount** and **mount** commands).

When you access the file with a data access or creation command, you can override some of the attributes that were set by a mount command or the server default settings.

Using mount to override server default attributes

In this example, the **mount** command is used to modify two processing attributes, specifying **binary** (rather than **text**), and **readtimeout(30)** (rather than the server default **readtimeout** value):

```
# mount mvshost1:"smith,binary,readtimeout(30)" /u/smith/mnt
```

In this example:

Operand	Description
mvshost1	Specifies the name of the host server.
smith	Specifies the name of the high-level qualifier of the MVS data sets.
binary	Specifies the processing mode for file contents sent between z/OS and the client (binary mode avoids ASCII/EBCDIC conversion).
readtimeout(30)	Specifies the amount of time (30 seconds) allowed since the last read access before the file is closed.
/u/smith/mnt	Specifies the name of the mount point (preferably an empty directory).

Getting authorization to access files

If the mount fails, check with your system administrator to ensure that you are authorized to access the MVS data sets or z/OS UNIX files and that the data sets or files are listed in the exports data set. The privilege level required to enter **mount** and **umount** commands varies among client operating system implementations. Many UNIX implementations limit these commands to the root user or superuser mode. If the MVS system operator issues the *freeze=on* operand of the **modify** command, all new tries to mount an MVS or z/OS UNIX file system fail until the z/OS system operator issues the *freeze=off* operand. If the z/OS system operator issues the *freeze=onhfs* operand of the modify command, conventional MVS data sets can still be mounted, but all new tries to mount z/OS UNIX file systems fail until the system operator issues the *freeze=offhfs* operand.

Saving of mount points

Once the **mount** command is issued successfully and a mount point is established between a local directory and the MVS or z/OS UNIX file system, the mount point information is saved in the mount handle data set by the z/OS NFS server. This information is used to automatically reestablish active mount points when the server is started. When accessing conventional MVS data sets, a mount point is active if the last activity against this mount point is less than the **restimeout** attribute value set by the system administrator.

The **mount** command does not need to be reissued for the same mount point in further sessions unless the **umount** command has been used to disconnect the mount point or the mount point has been disconnected by the cleanup activity of the **restimeout** site attribute. For more information about the **restimeout** site attribute see Chapter 9, “Initialization attributes for the z/OS NFS server,” on page 87.

Automatic timed logout - logout attribute

If there is no activity on the client within the period specified in the **logout** attribute of the attributes file, or the server stops, the connection between the server and the client workstation is logged out automatically. You must issue the **mvslogin** command again to get access to the z/OS files.

Displaying default and mount point attributes - showattr

Use the **showattr** command to display the default attributes or the attributes that have been set for a specific mount point. If you specify a mount point, showattr shows the attributes for the mount point, including the overriding values. For descriptions of the attributes, see Chapter 9, “Initialization attributes for the z/OS NFS server,” on page 87.

If you omit the hostname, you must specify the `/localpath`.

Table 8 shows examples of the showattr command for some clients.

Table 8. Examples of the showattr command for clients

Client Environments	Command Examples
AIX, UNIX	`showattr mvshost1 /u/smith/mnt`
z/OS	`showattr mvshost1 /u/smith/mnt`
Solaris	`showattr mvshost1 /u/smith/mnt`
Linux	`showattr mvshost1 /u/smith/mnt`

Make sure that your version of the showattr command matches the release of the z/OS NFS that you are using. Otherwise, the z/OS NFS server attributes do not display.

These examples show different ways you can use the **showattr** and **mount** commands.

Figure 5 shows a showattr command with just the host name (`mvshost1` in this example) specified. The default attributes for the server are displayed.

```
$ showattr mvshost1

z/OS NFS Server Data Set Creation Attributes:

lrecl(80)               recfm(fb)               blksize(3200)
space(10,10)            blks                    dsorg(ps)
dir(5)                  unit()                  volume(nfs003)
recordsize(512,4K)      keys(64,0)              nonspanned
shareoptions(1,3)
mgmtclas()              dsntype(pds)            norlse
dataclas()              storclas()

z/OS NFS Server Processing Attributes:

text                    lf                      blankstrip
nofastfilesize          retrieve                maplower
mapleaddot              executebitoff           setownerroot
attrtimeout(120)        readtimeout(90)         writetimeout(30)
sync                    nofileextmap            xlat(OEMVS)
srv_ccsid(0)            cln_ccsid(0)            nordrverf
sidefile(f066440.nfs.mapping)

z/OS NFS Server Site Attributes (not modifiable):

mintimeout(1)           nomaxtimeout            logout(1800)
nfstasks(8,8)           restimeout(48,0)        hfs(/hfs)
bufhigh(32M)            readaheadmax(16K)       cachewindow(128)
percentsteal(20)        maxrdforszleft(32)      logicalcache(16M)
smf(none)               nopcnfsd     security(exports,exports,exports)
leadswitch              sfmax(2)                nochecklist
fn_delimiter(,)         readdirtimeout(30)
public(IBMUSER.WEBNFS,/HFS/)
```

Figure 5. Displaying default attributes

If you use the terse (-t) option, the attributes display like this:

```
   $ showattr -t mvshost1
lrecl(8196),recfm(vb),blksize(0),space(100,10),trks,dsorg(ps),dir(27),unit(),
volume(vm5u16),recordsize(512,4K),keys(64,0),nonspanned,shareoptions(1,3),
mgmtclas(),dsntype(pds),rlse,dataclas(),storclas()
binary,lf,blankstrip,nofastfilesize,noretrieve,maplower,mapleaddot,
executebiton,setownerroot,attrtimeout(120),readtimeout(90),writetimeout(30),
sync,nofileextmap,xlat(oemvs),srv_ccsid(1047),cln_ccsid(819),notag,
convserv(lre),nordrverf,sidefile()
mintimeout(1),nomaxtimeout,logout(18000),nfstasks(8,8,8),restimeout(48,0),
hfs(/hfs),bufhigh(32M),readaheadmax(16K),cachewindow(112),percentsteal(20),
maxrdforszleft(32),logicalcache(16M),smf(none,off),pcnfsd,
security(none,none,none),leadswitch,sfmax(0),nochecklist,fn_delimiter(,),
readdirtimeout(30),hfsfbtimeout(30),upcase,rec878 ,mintasks(4,4,4),noremount,
nonlm,nodhcp,leasetime(120),public(),mvssec(sys,krb5,krb5i,krb5p),
hfssec(sys,krb5,krb5i,krb5p),pubsec(sys,krb5,krb5i,krb5p)
```

Figure 6 on page 56 illustrates the **showattr** command being used to display the attributes for the z/OS host named `mvshost1` as well as the mount point, `/u/smith/mnt`.

Figure 6 also illustrates the specified options. In the second showattr command, the client user specifies `/u/smith/mnt` in addition to `mvshost1`. This shows the user's specified settings at that mount point, rather than the default settings in the attributes data set.

```
# mount mvshost1:"smith,text,space(5,0),trks" /u/smith/mnt
$ showattr mvshost1

z/OS NFS Server Data Set Creation Attributes:

lrecl(8196)             recfm(vb)               blksize(0)
space(100,10)           trks                    dsorg(ps)
dir(27)                 unit()                  volume(vm5u16)
recordsize(512,4K)      keys(64,0)              nonspanned
shareoptions(1,3)
mgmtclas()              dsntype(pds)            rlse
dataclas()              storclas()

z/OS NFS Server Processing Attributes:

binary                  lf                      blankstrip
nofastfilesize          noretrieve              maplower
mapleaddot              executebiton            setownerroot
attrtimeout(120)                readtimeout(90)         writetimeout(30)
sync                    nofileextmap            xlat(oemvs)
srv_ccsid(1047)         cln_ccsid(819)          notag
convserv(lre)           nordrverf               sidefile()

z/OS NFS Server Site Attributes (not modifiable):

mintimeout(1)           nomaxtimeout            logout(18000)
nfstasks(8,8,8)         restimeout(48,0)        hfs(/hfs)
bufhigh(32M)            readaheadmax(16K)       cachewindow(112)
percentsteal(20)        maxrdforszleft(32)      logicalcache(16M)
smf(none,off)           pcnfsd                  security(none,none,none)
leadswitch              sfmax(0)                nochecklist
fn_delimiter(,)         readdirtimeout(30)      hfsfbtimeout(30)
upcase                  rec878                  mintasks(4,4,4)
noremount
nonlm                   nodhcp                  leasetime(120)
public()
mvssec(sys,krb5,krb5i,krb5p)
hfssec(sys,krb5,krb5i,krb5p)
pubsec(sys,krb5,krb5i,krb5p)
```

Figure 6. Displaying default and mount point attributes (Part 1 of 2). The client user changed the space(100,10), blks, and binary attributes to space(5,0), trks, and text for the mount point `/u/smith/mnt`, and then specified that mount point in the second showattr command.

```
$ showattr mvshost1 /u/smith/mnt

z/OS NFS Server Data Set Creation Attributes:

lrecl(8196)             recfm(vb)               blksize(0)
space(5,0)              trks                    dsorg(ps)
dir(27)                 unit()                  volume()
recordsize(512,4K)      keys(64,0)              nonspanned
shareoptions(1,3)
mgmtclas()              dsntype(pds)            norlse
dataclas()              storclas()

z/OS NFS Server Processing Attributes:

text                    lf                      blankstrip
nofastfilesize          retrieve                maplower
mapleaddot              executebitoff           setownerroot
attrtimeout(120)        readtimeout(90)         writetimeout(30)
sync                    nofileextmap            xlat(OEMVS)
srv_ccsid(0)            cln_ccsid(0)            nordrverf
sidefile(hlq.nfs.mapping)

z/OS NFS Server Site Attributes (not modifiable):

mintimeout(1)           nomaxtimeout            logout(1800)
nfstasks(8,8)           restimeout(48,0)        hfs(/hfs)
bufhigh(32M)            readaheadmax(16K)       cachewindow(128)
percentsteal(20)        maxrdforszleft(32)      logicalcache(16M)
smf(userfile)           nopcnfsd     security(exports,exports,exports)
leadswitch              sfmax(0)                checklist
fn_delimiter(,)         readdirtimeout(30)
public(NFSUSR3,/HFS/usr/lpp/internet/server_root/pub)
```

Figure 6. Displaying default and mount point attributes (Part 2 of 2). The client user changed the space(100,10), blks, and binary attributes to space(5,0), trks, and text for the mount point `/u/smith/mnt`, and then specified that mount point in the second showattr command.

Unmounting and logging out of z/OS

This section describes the **umount** and **mvslogout** commands.

Disconnecting your mount point - umount

Use the **umount** command to break the connection between the mount point on your client and the server. When you issue this client command, the file you were editing is released (written to DASD). You do not need to unmount after each session, unmount only when you no longer have a need to access the MVS file system. Check the documentation for your client operating system to ensure that you enter the umount command correctly.

Table 9 shows examples of the umount command for some clients.

Table 9. Examples of the umount command for clients

Client Environments	Umount Command Examples
AIX, UNIX	umount /u/smith/mnt
Solaris	umount /u/smith/mnt
Windows	nfs unlink z: (where z: is the NFS mounted drive)
Linux	umount /u/smith/mnt

In this example:

Operand	Description
u/smith/mnt	Specifies the mount point on the local file system.
mvshost1	Specifies the name of the z/OS host system.

For example, suppose that you want to unmount from the server, and the mount point on your workstation is named `/u/smith/mnt`. You could enter the **umount** command as follows:

```
# umount /u/smith/mnt
```

"No Such File or Directory" Message - The z/OS system operator can also unmount your workstation from the server. If this happens before you try to unmount, you get a "No such file or directory" error message.

Ending your z/OS session - mvslogout

Use the **mvslogout** command to disconnect from the z/OS host. The mvslogout command is only required when the **mvslogin** command was used to begin the connection.

An mvslogout to an z/OS user ID cancels a prior mvslogin to the same z/OS user ID from the same local host.

Your account is automatically logged out if it is inactive for the period of time specified in the **logout** site attribute.

Table 10 shows an example of the mvslogout command for some clients, in which the name of the z/OS host is `mvshost1`.

Table 10. Example of the mvslogout command for clients

Client Environments	Mvslogout Command Examples
AIX, UNIX	mvslogout mvshost1
Solaris	mvslogout mvshost1
Linux	mvslogout mvshost1

If you log out successfully, a message like this appears:

```
GFSA958I uid 215 logged out ok.
```

Chapter 7. Commands and examples for z/OS NFS clients

This chapter gives the syntax and examples of commands that z/OS NFS users need to know to access AIX, UNIX, z/OS, and other remote files using the z/OS NFS client. This chapter shows how to perform the following tasks.

- Log on to a remote z/OS system from the z/OS NFS client if the target server is a remote z/OS NFS server
- Access NFS files from the z/OS NFS client
- Display z/OS NFS client statistical data
- Query mount points
- Display default mount point attributes
- Mount and unmount remote file systems from the z/OS NFS client
- Log out of z/OS, if the target server was a remote z/OS NFS server

The command programs are intended to run in an z/OS shell environment and are not implemented as TSO/E commands, with the exception of **mount** and **unmount**.

The **mount** and **unmount** commands are not part of z/OS NFS. See *z/OS UNIX System Services Command Reference* for additional details. An example of the syntax and usage is shown here for your information. You can use the TSO HELP MOUNT and TSO HELP UNMOUNT commands to see the syntax that is applicable to your system.

For information about the z/OS NFS attributes see Chapter 8, "Initialization attributes for the z/OS NFS client," on page 79 and Chapter 9, "Initialization attributes for the z/OS NFS server," on page 87.

Using commands on DFSMS/MVS

The following is a summary of the syntax for the commands described in this chapter. See "Mount command syntax and examples" on page 65 and "Unmount command syntax and examples" on page 71 for information about the **mount** and **unmount** commands.

The **mvslogin** command is used to log in to z/OS from your workstation. The mvslogin command can be issued multiple times, and the last one overrides the previous one. The mvslogin command is only required when accessing data on systems where the z/OS NFS server site ***security*** attribute is set to *saf* or *safexp*.

The following is the mvslogin command syntax.

© Copyright IBM Corp. 1991, 2006

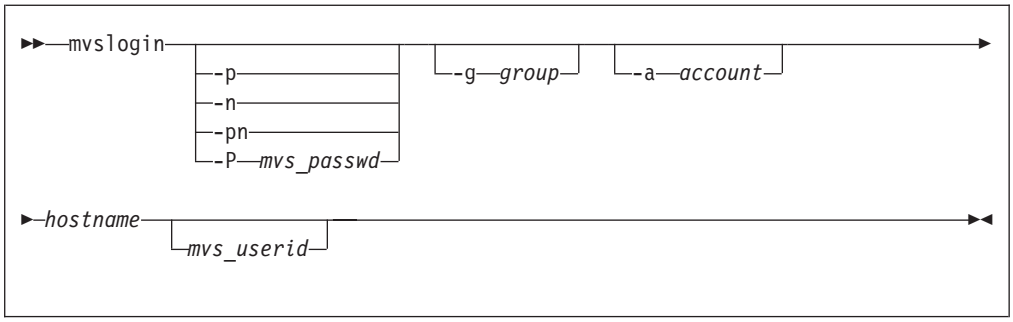

where

Operand	Description
-p	Causes a prompt for the user's z/OS password. The password is passed to z/OS to validate the user logging in. Your security procedures determine whether you should use this parameter.
-n	Causes a prompt for a new password.
-pn	Causes a prompt for the user's current password and then causes two prompts for the user's new password.
-P *mvs_passwd*	No prompt for your z/OS password; just type your z/OS password after the -P. This enables you to automate your MVS login.
-g *group*	A group name string passed to z/OS for accounting purposes. The maximum length is 8 characters.
-a *account*	An account string passed to z/OS for accounting purposes. The maximum length is 16 characters.
hostname	The name of the z/OS host (for example, `mvshost1`). The default is the local host.
mvs_userid	A user ID that z/OS recognizes as valid. If you do not specify this parameter, your workstation user name is used. The z/OS NFS server does not support the use of an alias user ID or a mixed case user ID with the mvslogin command.

The **showattr** command is used to display the default attributes or the attributes that have been set for a specific mount point. If you specify a mount point, showattr shows the attributes for the mount point, including the overriding values. For descriptions of the attributes, see Chapter 9, "Initialization attributes for the z/OS NFS server," on page 87.

The following is the showattr command syntax.

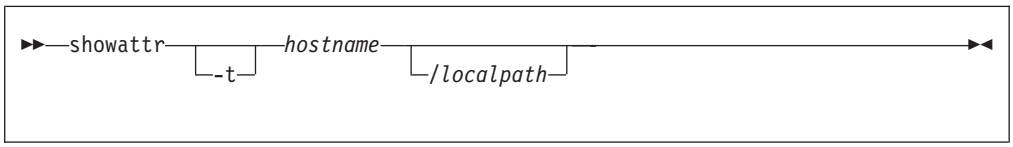

where

Operand	Description
-t	Used to specify tersed output.

hostname The name of the z/OS host (for example, `mvshost1`). The default is the local host.

/localpath The mount point on your client system (for example, `/u/smith/mnt`). This should be an empty directory.

The **nfsstat** command is used to display the z/OS NFS client statistical information, to reset the statistical information to zero, to display NFS mount point information, or to set the debug status.

The following is the nfsstat command syntax.

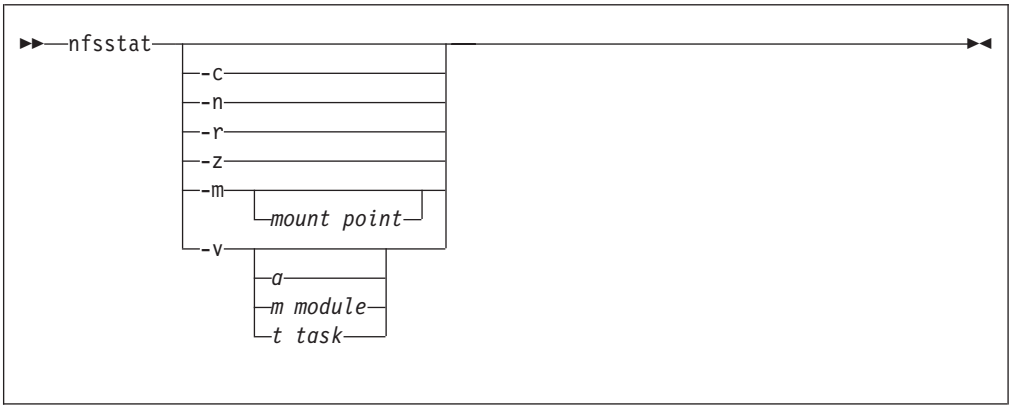

where

Operand	Description
-p	Causes a prompt for the user's z/OS password. The password is passed to z/OS to validate the user logging in. Your security procedures determine whether you should use this parameter.
-c	Displays both NFS and RPC statistics about the z/OS NFS client. This is the default option on the nfsstat command.
-n	Displays NFS statistics about the z/OS NFS client.
-r	Displays RPC statistics about the z/OS NFS client.
-z	Initializes statistics to zero. Used by root user only. This option can be combined with options *-c*, *-n*, and *-r* on the nfsstat command. When combined with these nfsstat options, each particular set of statistics is set to zero after the statistics are printed.
-m	Displays the name of each NFS mounted file system.
-m *mount point*	Displays information for the NFS mounted file system on the specified mount point.
-v	Returns information about the latest APAR installed on the z/OS NFS client.
-v a	Returns a list of all the modules in the z/OS NFS client with their current level information. At the end of the list is information about the latest APAR installed.
-v m *module*	Returns information about the APAR level of the specified *module*.
-v t *task*	Returns information about the APAR level of the specified *task*.

The **showmount** command is used to display the remote NFS server mount information. If you omit the options, the default option displays hostnames of all remote mounts from the *hostname* NFS server. If you omit the *hostname* parameter, then the local hostname is used.

The following is the showmount command syntax.

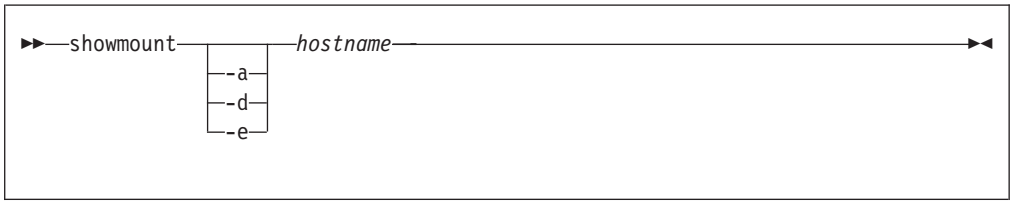

If you omit the options, the default option displays hostnames of all remote mounts from the *hostname* NFS server.

where

Operand	Description
-a	Displays all mounts in the format hostname:directory from the hostname specified in the showmount command.
-d	Displays only directory names of all mounts from the hostname specified in the showmount command.
-e	Displays the list of exported directories from the hostname specified in the showmount command.
hostname	The name of the NFS server host (for example, `mvshost1`). The default is the local host.

The **crnl2nl** command will remove the carriage-return and end-of-file mark (EOF) from the input file. No other conversions are performed. Either, or both the input and output file can be a local or remote file.

The following is the crnl2nl command syntax.

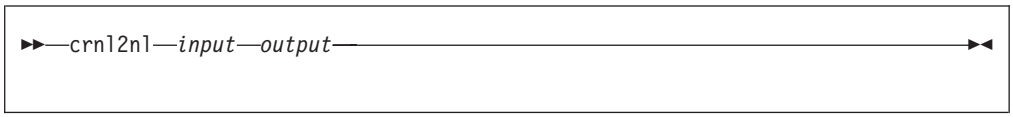

The data is assumed to be in EBCDIC format.

where

Operand	Description
input	Absolute path name of the input file to be converted.
output	Absolute path name of the output file.

The **nl2crnl** converts the line delimiter from MVS format to carriage-return newline (CRNL) format; the newline (NL) in the input file is converted to the carriage-return newline pairs (CR, NL). No other conversions are performed. Both input and/or output file can be local or remote file.

The following is the nl2crnl command syntax.

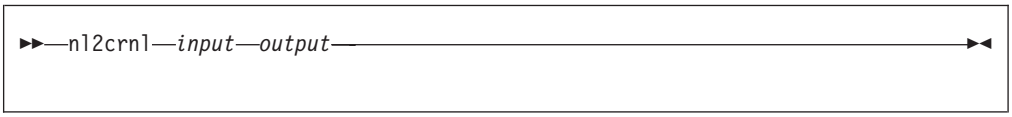

where

Operand	Description
input	Absolute path name of the input file to be converted.
output	Absolute path name of the output file.

Notes:

1. The size of the write buffer is double the size of the read buffer.
2. The data is assumed to be in EBCDIC format.

The **mvslogout** command is used to disconnect from the remote z/OS NFS server host. The mvslogout command is only required when the mvslogin command was used to begin the connection.

The following is the mvslogout command syntax.

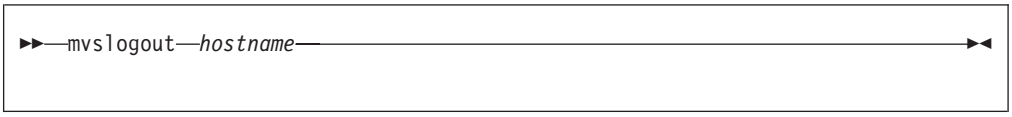

where

Operand	Description
hostname	The name of the z/OS host (for example, `mvshost1`). The default is the local host.

Accessing z/OS

To access remote z/OS files, enter both the **mvslogin** command to log in to the z/OS NFS server's host system and the **mount** command to mount the files or data sets to your local system. The mvslogin command is only required when accessing data on systems where the z/OS NFS server site *security* attribute is set to saf or safexp. Once the files or data sets are mounted to a local mount point, you can read, write, create, delete and treat the mounted files as part of your local z/OS UNIX files. When you are finished with your work, use the **unmount** and **mvslogout** commands to break the connection. The mvslogout command is only required when the mvslogin command was used to begin the connection.

Note: Issuing the mount and unmount commands, as well as creation of the mount point, must be performed by someone with superuser authority.

To access files on z/OS systems where the z/OS NFS server site *security* attribute is set to saf or safexp, you need a z/OS user ID and password, and authorization to access the files that you need. You can only establish one z/OS session for each z/OS user ID. If you do not already have an z/OS user ID, a z/OS password, and access authorization on the z/OS system from which you require NFS services, request them from your z/OS system administrator.

Note: If you cannot use the mvslogin, mvslogout, or **showattr** commands, they might be installed incorrectly. Ensure that your system administrator has made the executable code for these three commands available to your z/OS user ID and that you have been given the correct access authority to them.

Mvslogin command examples

Use the **mvslogin** command to log in to the remote z/OS system. The mvslogin command can be issued multiple times and the last one overrides the previous one. The mvslogin command is only required when accessing data on systems where the z/OS NFS server site *security* attribute is set to *saf* or *safexp*.

Table 11 shows examples of the mvslogin command where `mvshost1` is the name of the z/OS host and `smith` is the user's ID on z/OS.

Table 11. Examples of the mvslogin command for clients

Command Examples
mvslogin -p mvshost1 smith
mvslogin -P smithspw -g finance -a 5278 mvshost1 smith
mvslogin -n mvshost1 smith
mvslogin -pn -a 5278 mvshost1 smith
mvslogin mvshost1
mvslogin mvshost1 smith

In the example where the user enters **mvslogin mvshost1**, the current login client user ID is used as the z/OS user ID.

In the example where the user enters **mvslogin mvshost1 smith**, the system then prompts for Smith's z/OS password. If Smith logs in successfully, this message displays.

```
GFSA955I smith logged in ok.
```

Otherwise, an appropriate error message displays.

Note: Messages that start with GFSA and GFSC apply to NFS requests. See Chapter 19, "Network File System messages," on page 241.

"Permission denied" message

If you have successfully logged in and get the "Permission denied" message while trying to access data, that can be due to one of these cases:

- An **mvslogout** command for the same z/OS host has been entered from the same client platform. See mvslogout for details.
- Your z/OS user ID has been automatically logged out because the **logout** attribute value has been exceeded. This can happen when you leave the client idle for too long. Enter the **mvslogin** command again, and start your processes again. To find out how many seconds you can stay logged in while your client is idle, issue the **showattr** command and look at the logout attribute.
- Another mvslogin to the same z/OS host using the same z/OS ID has been entered from the same UID and the same client platform. If this is the case, retry your access.

Mount command syntax and examples

Use the TSO MOUNT command to make a connection between a mount point on your local file system and one or more files on a remote AIX, UNIX, z/OS, or other file system. The MOUNT command can only be used by a z/OS superuser.

Note: The same mount function can also be performed using the z/OS UNIX automount facility or /etc/rc shell scripts support. z/OS UNIX does not support NFS mounts in the SYS1.PARMLIB member statement. When the automount facility is used to manage remote NFS mount points, the z/OS NFS client user could experience ESTALE/EIO errors if the automounter unmounts the accessed mount point when the time limits specified by the automount DURATION and DELAY parameters have been exceeded. The automount default, DURATION=NOLIMIT, disables the DURATION timeout period. The DELAY for unmounting file systems should be longer than the time z/OS NFS clients are likely to keep NFS mounts to the files and directories active. For more information see *z/OS UNIX System Services File System Interface Reference*.

Figure 7 illustrates the syntax of the TSO MOUNT command. For more information about the **mount** command, see *z/OS UNIX System Services Command Reference*.

```
MOUNT FILESYSTEM(file_system_name)
      TYPE(NFS)
      MOUNTPOINT(local_mountpoint)
      MODE(RDWR|READ)
      PARM('hostname:"path_name,server_attributes", client_options')
      SETUID|NOSETUID
      TAG(TEXT,CCSID)
      WAIT|NOWAIT
```

Figure 7. TSO MOUNT command syntax operands

where

FILESYSTEM(*file_system_name*)
Specifies the name of the file system to be added to the file system hierarchy. This operand is required. The file system name specified must be unique among previously mounted file systems. It may be an arbitrary name up to 44 characters in length of a filesystem. You can enclose file_system_name in single quotes, but they are not required.

TYPE(*NFS*)
Specifies the type of file system that performs the logical mount request. The NFS parameter must be used.

MOUNTPOINT(*local_mountpoint*)
Specifies the path name of the mount point directory, the place within the file hierarchy where the file system is to be mounted. The local_mountpoint specifies the mount point path name. The name can be a relative path name or an absolute path name. The relative path name is relative to the current working directory of the TSO/E session (usually the home directory). Therefore, you should usually specify an absolute path name. A path name is case-sensitive, so enter it exactly as it is to appear. Enclose the path name in single quotes.

Note: The mount point must be a directory. Any files in that directory are inaccessible while the file system is mounted. Only one file system can be mounted to a mount point at any time.

MODE(*RDWR*|*READ*)
Specifies the type of access for which the file system is to be opened.

RDWR specifies that the file system is to be mounted for read and write access. *RDWR* is the default option.

READ specifies that the file system is to be mounted for read-only access.

PARM(*'hostname:"path_name,server_attributes", client_options'*)
Specifies the hostname of the remote Network File System server, the server attributes, and the options. See Chapter 9, "Initialization attributes for the z/OS NFS server," on page 87 for descriptions of the server attributes, including the **mvsmnt** attribute to be used with mount commands using the NFS version 4 protocol.

The double quotes can be omitted if no server attributes are specified. Enclose the entire string in single quotes. If the automount facility is being used, the single quotes may be specified but are not required. You can specify lowercase or uppercase characters. A blank space must preceed any client_options specified.

SETUID|**NOSETUID**
See *z/OS UNIX System Services Command Reference* for details about **SETUID** and **NOSETUID**.

TAG(*TEXT,CCSID*)
See *z/OS UNIX System Services Command Reference* for details about the **TAG** keyword.

Note: When TAG is specified, *xlat(Y)* must not be specified; otherwise, mount will fail.

WAIT|**NOWAIT**
See *z/OS UNIX System Services Command Reference* for details on **WAIT** and **NOWAIT**.

Data conversion

The z/OS NFS client supports data conversion defined by the universal character encoding standard known as the Unicode Standard on z/OS V1R2 (and above) when reading data from a remote NFS server or writing data to a remote z/OS NFS server. The Unicode Standard offers character conversion as well as basic case conversion. Within character conversion, characters are converted from one coded character set identifier (CCSID) to another. CCSID information is obtained from the *cln_ccsid* and *srv_ccsid* parameters.

On z/OS releases below V1R2, the z/OS NFS client supports data conversion defined by the Character Data Representation Architecture (CDRA) when reading data from a remote NFS server or writing data to a remote NFS server. CDRA characters are converted from one CCSID to another. CCSID information is obtained from the cln_ccsid and srv_ccsid parameters. See *Character Data Representation Architecture Reference and Registry* and *Character Data Representation Architecture Overview* for additional information.

Only single byte to single byte data conversion is supported. For example, if a client file has a CCSID of 437 and a server file has a CCSID of 297, data conversion will occur between USA ASCII format (CCSID 437) and French EBCDIC format (CCSID 297). Single byte to multiple byte conversion (including double byte character set (DBCS)) is not supported and will result in NFS terminating with an error message.

The **cln_ccsid**, **srv_ccsid**, and **xlat** parameters identify whether data conversion is performed, and how data conversion is done. The **cln_ccsid**, **srv_ccsid**, and **xlat** parameters are supported by the z/OS NFS client installation parameter and TSO MOUNT command. The parameters on a TSO MOUNT command override the parameters specified as a z/OS NFS client installation parameter.

BSAM, QSAM, and VSAM ESDS access to remote files

BSAM, QSAM, and VSAM ESDS applications can access files stored on remote NFS servers through the z/OS NFS client. This will allow existing MVS application programs access to data on other systems using BSAM, QSAM, and VSAM ESDS interfaces. The BSAM, QSAM, and VSAM ESDS access methods assume that all text files are EBCDIC. When using these access methods, the **delim** parameter indicates whether the remote files contain text or binary data. Text data consists of records that are separated by a delimiter. If the delim parameter is not binary, the EBCDIC text delimiter is used by the access methods when processing the remote files. The **delim** parameter is supported on the z/OS NFS client installation parameter and TSO MOUNT command.

Note: The z/OS NFS client can also access VSAM key-sequenced (KSDS) and relative record (RRDS) data sets.

All the remote file objects under the same mount point have the same **delim** value. The **delim** parameter cannot be set on a file basis under the mount point. The **delim** parameter in the TSO MOUNT command overrides the **delim** parameter specified in z/OS NFS client installation parameter. However, you can override the **delim** parameter on the TSO MOUNT command with the **filedata** parameter on a JCL DD statement, SVC 99, or TSO ALLOCATE command. The **filedata** parameter can be either **text** or **binary**.

For BSAM, QSAM, and VSAM ESDS applications accessing files stored on remote NFS servers, the z/OS NFS client provides data conversion when the *xlat=Y* parameter is specified. It uses the cln_ccsid and srv_ccsid settings. When *xlat=N*, the z/OS NFS client will not perform data conversion. The filedata parameter on a JCL DD statement is also used to specify if the data consists of text records separated by delimiters or if the data is binary and does not contain record delimiters. To avoid undesirable data conversions, care should be taken to insure the specification of the xlat and delim parameters are not in conflict with the data type specified by the filedata parameter on a JCL DD statement. The filedata and delim parameters only affect BSAM, QSAM, and VSAM ESDS data access and have no affect on the z/OS NFS client data conversion. The z/OS NFS client data conversion is only controlled by the xlat parameter. The significance of different filedata and delim combinations are described in the following information.

Note: In each case insure the z/OS NFS client *xlat=Y*, *cln_ccsid*, and *srv_ccsid* parameter settings are correct for the filedata and delim combination.

FILEDATA=TEXT, *delim*=notBINARY
: Specifies that the data is to be accessed as text. The access method appends a record delimiter on output and expects delimiters on input. The delimiter used is that specified on the ***delim*** parameter.

FILEDATA=TEXT, *delim*=BINARY
: Specifies that the data is to be accessed as text. The access method appends a record delimiter on output and expects delimiters on input. The delimiter used is the default of the EBCDIC new line character (x'15') since the delim parameter does not specify a valid text delimiter.

FILEDATA=BINARY, *delim*=notBINARY
Specifies that the data is to be accessed as binary. The access method does not append record delimiters on output, does not recognize record delimiters on input, and it treats all characters as data on input.

FILEDATA=BINARY, *delim*=BINARY
Specifies that the data is to be accessed as binary. The access method does not append record delimiters on output, does not recognize record delimiters on input, and it treats all characters as data on input.

FILEDATA=not specified, *delim*=specified
Means that the data is to be accessed according to the value specified in the delim parameter.

FILEDATA=not specified, *delim*=not specified
Means that the data is to be accessed as binary. The access method does not append record delimiters on output, does not recognize record delimiters on input, and it treats all characters as data on input.

The z/OS NFS client also provides UNIX authentication for security and provides the z/OS UNIX client's user ID (UID), group ID (GID), and a list of users GIDs to the remote NFS server for authorization checking. When the remote NFS server is the z/OS NFS server, the **mvslogin** command can be used to provide additional security checking through RACF authentication. MVS application programs which require access to data on remote z/OS systems may be required to perform an mvslogin.

For information on BSAM, QSAM, and VSAM ESDS applications access to z/OS UNIX or remote files and their restrictions, refer to *z/OS DFSMS Using Data Sets*.

Mount command examples

Use the TSO MOUNT command to make a connection between a local mount point on the z/OS NFS client and an NFS server.

In this example, the **mount** command is used to mount a set of MVS files. The PARM operand contains the z/OS NFS server **text** processing attribute and requires the use of double quotes around the string user,text.

```
mount filesystem(nfs00) type(nfs) mountpoint('/u/nfsdir')
parm('stlmvs3:"user,text",soft,timeo(100)')
```

In this example:

Operand	Description
nfs00	Specifies the name of the file system to be added to the file system hierarchy.
nfs	Specifies the required file system type.
/u/nfsdir	Specifies the name of the mount point (preferably an empty directory).
stlmvs3	Specifies the name of the host z/OS NFS server.
user	Specifies the name of the high-level qualifier of the MVS files on the z/OS NFS server.
text	Specifies the processing attribute for the z/OS NFS server.
soft	Specifies the `PARM` operand option for the z/OS NFS client.
timeo(100)	Specifies the `PARM` operand option for the z/OS NFS client.

In this example, the mount command is used to mount a z/OS UNIX directory. The PARM operand contains the z/OS NFS server *binary* processing attribute and requires the use of double quotes around the string */hfs/u/user,binary*.

```
mount filesystem(nfs01) type(nfs) mountpoint('/u/nfsdir1')
parm('stlmvs3:"/hfs/u/user,binary",soft')
```

In this example:

Operand	Description
nfs01	Specifies the name of the file system to be added to the file system hierarchy.
nfs	Specifies the required file system type.
/u/nfsdir1	Specifies the name of the mount point (preferably an empty directory).
stlmvs3	Specifies the name of the host z/OS NFS server.
/hfs/u/user	Specifies the name of the z/OS UNIX directory on the z/OS NFS server.
binary	Specifies the processing attribute for the z/OS NFS server.
soft	Specifies the PARM operand option for the z/OS NFS client.

In this example, the mount command is used to mount an AIX home directory.

```
mount filesystem(nfs02) type(nfs) mountpoint('/u/nfsdir2')
parm('aix6000:/home/user,xlat(y)')
```

In this example:

Operand	Description
nfs02	Specifies the name of the file system to be added to the file system hierarchy.
nfs	Specifies the required file system type.
/u/nfsdir2	Specifies the name of the mount point (preferably an empty directory).
aix6000	Specifies the name of the host AIX NFS server.
/home/user	Specifies the name of the home directory on the AIX NFS server.
xlat(y)	Specifies the PARM operand option for the z/OS NFS client.

In this example, the mount command is used to mount an AIX home directory using the NFS version 3 protocol.

```
mount filesystem(nfs03) type(nfs) mountpoint('/u/nfsdir2')
parm('aix6000:/home/user,xlat(y),vers(3)')
```

In this example:

Operand	Description
nfs03	Specifies the name of the file system to be added to the file system hierarchy.
nfs	Specifies the required file system type.
/u/nfsdir2	Specifies the name of the mount point (preferably an empty directory).

aix6000	Specifies the name of the host AIX NFS server.
/home/user	Specifies the name of the home directory on the AIX NFS server.
xlat(y)	Specifies the PARM operand option for the z/OS NFS client.
vers(3)	Specifies the version of NFS protocol that is being used.

Automount facility mount: In this example, the mount command in the automount policy file is used to do ASCII to EBCDIC conversion.

```
name  *
type  NFS
filesystem shared.<asis_name>
mode  rdwr
duration 15
delay 10
parm  mvshost:"/export/<asis_name>", xlat(Y)
setuid no
```

In this example:

Operand	Description
type NFS	Identifies the automount as NFS.

For more information about the z/OS UNIX automount facility, see *z/OS UNIX System Services Planning* and *z/OS UNIX System Services File System Interface Reference*.

Getting authorization to access files

If the mount fails, check with your system administrator to ensure that you are authorized to access the AIX or UNIX file system listed in the exports control file. The z/OS NFS client also provides UNIX authentication for security and provides the z/OS UNIX client's UID, GID, and a list of the GIDs from the z/OS UNIX client's groups to the remote NFS server for authorization checking. When the remote NFS server is the z/OS NFS server, the **mvslogin** command can be used to provide additional security checking through RACF authentication. The privilege level required to enter **mount** and **unmount** commands is superuser. When requesting service from the z/OS NFS server, if the z/OS system operator issues the *freeze=on* operand of the **modify** command, all new tries to mount a z/OS UNIX file system fail until the z/OS system operator issues the *freeze=off* operand. If the z/OS UNIX system operator issues the *freeze=onhfs* operand of the modify command, conventional MVS data sets can still be mounted, but all new tries to mount z/OS UNIX files fail until the system operator issues the *freeze=offhfs* operand.

Saving of mount points

Once the **mount** command is issued successfully and a mount point is established between a remote directory and the file system, the mount point information is not saved by the z/OS NFS client. The z/OS NFS client does not maintain mount persistence across restart. If z/OS UNIX or the z/OS NFS client is restarted, all prior session's mount point information is lost and all mount points must be reestablished.

Automatic timed logout - logout attribute

When using Network File System services from the z/OS NFS server, if there is no activity on the client within the period specified in the **logout** attribute of the attributes file, or if the server stops, the connection between the server and the client workstation is logged out automatically. You must issue the **mvslogin** command again to get access to the z/OS files.

Unmount command syntax and examples

This section describes the **unmount** command.

Disconnecting your mount point - unmount

Use the **unmount** command to break the connection between the mount point on your client and the server (that is, to unmount). You must have superuser authority to issue the unmount command.

Note: The same unmount function can also be performed using the z/OS UNIX **automount** facility or /etc/rc shell scripts support. When the automount facility is used to manage remote NFS mount points, the z/OS NFS client user could experience ESTALE/EIO errors if the automounter unmounts the accessed mount point when the time limits specified by the automount ***duration*** and ***delay*** parameters have been exceeded. For additional information on the z/OS UNIX automount facility (**/etc/rc** shell scripts support) see *z/OS UNIX System Services Planning* and *z/OS UNIX System Services File System Interface Reference*.

Figure 8 illustrates the syntax of the TSO UNMOUNT command. For more information about the UNMOUNT command, see *z/OS UNIX System Services Command Reference*.

```
UNMOUNT  FILESYSTEM(file_system_name)
NORMAL | DRAIN | IMMEDIATE | FORCE | RESET
```

Figure 8. TSO UNMOUNT command syntax operands

where

FILESYSTEM(*file_system_name*)
Specifies the name of the file system to be removed from the file system hierarchy. file_system_name specifies the file_system_name exactly as it was specified when the file system was originally mounted. You can enclose file_system_name in single quotes, but they are not required.

NORMAL | DRAIN | IMMEDIATE | FORCE | RESET
NORMAL: Specifies that if no user is accessing any of the files in the specified file system, the unmount request is processed. Otherwise, the system rejects the unmount request. **NORMAL** is the default option.

DRAIN: Specifies that the system is to wait until all uses of the file system have ended normally before the unmount request is processed or until another UNMOUNT command is issued.

Note: UNMOUNT can be specified with IMMEDIATE to override a previous unmount DRAIN request for a file system. If this is used in the foreground, your TSO/E session waits until the unmount request has completed. The attention request key (usually ATTN or PA1) will not end the command.

IMMEDIATE: Specifies that the system is to unmount the file system immediately. Any users accessing files in the specified system receive failing return codes. All data changes to files in the specified file system are saved. If the data changes to files cannot be saved, the unmount request fails.

Note: UNMOUNT of an NFS mount point (regardless of soft or hard mount option) with NORMAL, DRAIN, or IMMEDIATE may fail with the return code of EBUSY if the z/OS NFS client determines that there are ongoing NFS requests to the NFS server. The UNMOUNT receiving EBUSY can have the file system unmounted immediately with FORCE, at the risk of data loss.

FORCE: Specifies that the system is to unmount the file system immediately. Any users accessing files in the specified file system receive failing return codes. All data changes to files in the specified file are saved, if possible. If the data changes cannot be saved to the files, the unmount request continues and data is lost.

Note: You must issue an UNMOUNT IMMEDIATE request before issuing UNMOUNT FORCE. Otherwise, UNMOUNT FORCE fails.

RESET: A reset request stops a previous unmount DRAIN request.

The following example unmounts the file system NFSC_001 normally:

```
UNMOUNT FILESYSTEM('NFSC_001')
```

The following example forces an **unmount** of the file system NFSC_001. You must issue an UNMOUNT IMMEDIATE before you can issue an unmount *FORCE* command.

```
UNMOUNT FILESYSTEM('NFSC_001') IMMEDIATE
UNMOUNT FILESYSTEM('NFSC_001') FORCE
```

If you receive a **"No Such File or Directory"** message, the z/OS system operator can also unmount your workstation from the server. If this happens before you try to unmount, you get a "No such file or directory" error message.

Displaying client statistical information–nfsstat

Use the **nfsstat** command to display the z/OS NFS client statistical information, to reset the statistical information to zero, to display NFS mount point information, or to set the debug status.

The following nfsstat command displays the NFS and RPC statistics for the z/OS NFS client:

```
nfsstat -c
```

The following nfsstat command initializes the NFS and RPC statistics for the z/OS NFS client to zero. This option may be used by the root user only:

```
nfsstat -z
```

Figure 9 on page 73 shows the output from the nfsstat command using the -c option to display the RPC and NFS statistics for the NFS client.

```
USER1:/u/user1:>nfsstat -c

Client rpc:
calls           badcalls        retrans         timeout
51              0               0               0

Client nfs_V2:
calls        badcalls
45           0
null         getattr     setattr          root         lookup
0        0%  21     47%  1           2%   0        0%  13         29%
readlink     read        writecache       write        create
0        0%  0       0%  0           0%   0        0%  1           2%
remove       rename      link             symlink      mkdir
4        9%  0       0%  0           0%   0        0%  0           0%
rmdir        readdir     fsstat
0        0%  3       7%  2           4%


Client nfs_V3:

calls           badcalls
4            0
null         getattr      setattr          lookup       access
0        0%  0        0%  0           0%   0        0%  0        0%
readlink     read         write            create       mkdir
0        0%  0        0%  0           0%   0        0%  0        0%
symlink      mknod        remove           rmdir        rename
0        0%  0        0%  0           0%   0        0%  0        0%
link         readdir      readdirplus      fsstat       fsinfo
0        0%  0        0%  0           0%   2        50%  1       25%
pathconf     commit
1       25%  0        0%
```

Figure 9. Displaying NFS client rpc and NFS statistical information

Figure 10 shows the output from the nfsstat command using the -r option to display the remote procedure call (RPC) statistics for the NFS client.

```
#nfsstat -r
Client rpc:
calls           badcalls        retrans         timeout
61              0               20005           20005
```

Figure 10. Displaying NFS client rpc statistical information

In this example:

Operand	Description
calls	Specifies the total number of RPC calls sent.
badcalls	Specifies the total of RPC calls rejected by a server.
retrans	Specifies the number of times an RPC call had to be retransmitted.
timeout	Specifies the number of times an RPC call timed out.

Figure 11 on page 74 shows the output from the **nfsstat** command using the -n option to display the NFS statistics for the NFS client.

```
# nfsstat -n

Client nfs_V2:
calls        badcalls
45           0
null         getattr     setattr         root         lookup
0        0%  21     47%  1          2%   0        0%  13         29%
readlink     read        writecache      write        create
0        0%  0       0%  0          0%   0        0%  1           2%
remove       rename      link            symlink      mkdir
4        9%  0       0%  0          0%   0        0%  0           0%
rmdir        readdir     fsstat
0        0%  3       7%  2          4%


Client nfs_V3:
calls           badcalls
4               0
null         getattr      setattr          lookup       access
0        0%  0        0%  0           0%   0        0%  0        0%
readlink     read         write            create       mkdir
0        0%  0        0%  0           0%   0        0%  0        0%
symlink      mknod        remove           rmdir        rename
0        0%  0        0%  0           0%   0        0%  0        0%
link         readdir      readdirplus      fsstat       fsinfo
0        0%  0        0%  0           0%   2        50%  1       25%
pathconf     commit
1       25%  0        0%
```

Figure 11. Displaying NFS client NFS statistical information

In this example:

Operand	Description
calls	Specifies the total number of NFS calls sent.
badcalls	Specifies the total of NFS calls rejected by a server.

Figure 12 shows the output from the nfsstat command using the -m option to display the server and path name of each NFS mounted file system.

```
#nfsstat -m
mvshost:/sj/sjpl is mounted on /sj/sjpl/host1
        filesystem NFS_MNT1
mvshost2:/sj/sjpl2,text mounted on /sj/sjpl/host2
        filesystem NFS_MNT2
```

Figure 12. Displaying NFS mounted file system information

Figure 13 and Figure 14 on page 75 show the output from the nfsstat command using the -m option to display the server name, path name, and attributes of mount point */sj/sjpl/host1* from both version 2 and version 3 protocols.

```
#nfsstat -m /sj/sjpl/host1
server  mvshost
path    /sj/sjpl/host1

hard    vers(2)   retry(10)   timeo(7)        retrans(3)
delim(binary)   xlat(n)       cln_ccsid(1047)  srvccsid(819)
rsize(8192)     wsize(8192)   readahead(8)    delaywrite(16)
acregmin(3)     acregmax(60)  acdirmin(30)    acdirmax(60)
datacaching(y)  attrcaching(y)  dynamicsizeadj(y)
```

Figure 13. Displaying V.2 protocol NFS mounted file system information

```
#nfsstat -m /sj/sjpl/host1
server  mvshost
path    /sj/sjpl/host1

hard    vers(3)   retry(10)   timeo(7)        retrans(3)
delim(binary)   xlat(n)       cln_ccsid(1047)  srvccsid(819)
rsize(32768)     wsize(32768)   readahead(8)    delaywrite(16)
acregmin(3)     acregmax(60)  acdirmin(30)    acdirmax(60)
datacaching(y)  attrcaching(y)  dynamicsizeadj(y)
```

Figure 14. Displaying version 3 protocol NFS mounted file system information

Displaying server mount information–showmount

Use the **showmount** command to display the remote NFS server mount information. If you omit the options, the default option displays hostnames of all remote mounts from the `hostname` NFS server. If you omit the *hostname* parameter, then the local hostname is used.

The following showmount command displays all remote mounts in the format `hostname:directory` from the local `hostname` NFS server.

```
showmount -a
```

The following showmount command displays only the directory names of all the remote mounts from the local *hostname* NFS server.

```
showmount -d
```

The following example shows the output from the showmount command using the -a option to display all mounts in the format hostname:directory from the hostname `mvshost`.

```
# showmount -a mvshost
mvshost.sanjose.ibm.com:/IBMUSER
usera.sanjose.ibm.com:/USER2
```

The following example shows the output from the showmount command using the -d option to display only the directory names of all mounts from the hostname `mvshost`.

```
# showmount -d mvshost
/IBMUSER
/USER2
```

The following example shows the output from the showmount command with no option specified to display only the hostnames of all remote mounts from the hostname `mvshost`.

```
# showmount mvshost
mvshost.sanjose.ibm.com
usera.sanjose.ibm.com
```

The following example shows the output from the showmount command using the -e option to display the exported directories from the hostname `aix_server1`.

```
USER1:/u/user1:>showmount -e aix_server1
Export list for host aix_server1:
/home/u/guest/test (everyone)
/usr/lpp/info      (everyone)
/tmp               (everyone)
```

The following example shows the output from the showmount command using the -e option to display the exported directories from the hostname `mvshost`. In this case, `mvshost` has the site attribute set to ***security***(*none*).

```
# showmount -e mvshost
No exported file systems for host MVSHOST
```

The following examples shows the output from the showmount command using the -e option to display the exported directories from the hostname `mvshost`. In this case, `mvshost` has the site attribute set to ***security***(*safexp*).

```
# showmount -e mvshost
Export list for host MVSHOST:
/IBMUSER              user1
```

Displaying default and mount point attributes–showattr

Use the **showattr** command to display the default attributes or the attributes that have been set for a specific mount point of the z/OS NFS server. If you specify a mount point, showattr shows the attributes for the mount point, including the overriding values. For descriptions of the attributes, see Chapter 9, "Initialization attributes for the z/OS NFS server," on page 87 and "Mount command syntax and examples" on page 65.

If you omit the *hostname*, you must specify the `/localpath`.

The following is an example of the showattr command.

```
showattr mvshost1 /u/smith/mnt
```

Make sure that your version of the showattr command matches the release of the z/OS NFS that you are using. Otherwise, the z/OS NFS server attributes will not display.

These examples show different ways you can use the showattr command.

Figure 15 on page 77 shows a showattr command with just the hostname (`mvshost1` in this example) specified. The default attributes for the server are displayed.

```
USER1:/u/user1:>showattr mvshost1

z/OS Network File System Server Data Set Creation Attributes:

lrecl(8196)             recfm(vb)               blksize(0)
space(100,10)           blks                    dsorg(ps)
dir(27)                 unit()                  volume()
recordsize(512,4K)      keys(64,0)              nonspanned
shareoptions(1,3)
mgmtclas()              dsntype(pds)            norlse
dataclas()              storclas()

z/OS Network File System Server Processing Attributes:

binary                  lf                      blankstrip
nofastfilesize          retrieve                maplower
mapleaddot              executebitoff           setownerroot
attrtimeout(120)        readtimeout(90)         writetimeout(30)
sync                    nofileextmap            xlat(OEMVS)
srv_ccsid(0)            cln_ccsid(0)            nordrverf
sidefile(hlq.nfs.mapping)

z/OS Network File System Server Site Attributes (not modifiable):

mintimeout(1)           nomaxtimeout            logout(1800)
nfstasks(8,8)           restimeout(48,0)        hfs(/hfs)
bufhigh(32M)            readaheadmax(16K)       cachewindow(128)
percentsteal(20)        maxrdforszleft(32)      logicalcache(16M)
smf(userfile)           pcnfsd                  security(exports,exports,exports)
leadswitch              sfmax(0)                nochecklist
fn_delimiter(,)         readdirtimeout(30)
public(NFSUSR3,/HFS/usr/lpp/internet/server_root/pub)
```

Figure 15. Displaying default attributes

If you use the terse (-t) option, the following attributes display.

```
USER1:/u/user1:>showattr -t mvshost1

lrecl(8196),recfm(vb),blksize(0),space(100,10),blks,dsorg(ps),dir(27),unit(),
volume(),recordsize(512,4K),keys(64,0),nonspanned,shareoptions(1,3),mgmtclas(),
dsntype(pds),norlse,dataclas(),storclas()
binary,lf,blankstrip,nofastfilesize,retrieve,maplower,mapleaddot,executebitoff,
setownerroot,attrtimeout(120),readtimeout(90),writetimeout(30),sync,
nofileextmap,xlat(oemvs),srv_ccsid(0),cln_ccsid(0),nordrverf,sidefile(hlq.nfs.mapping)
mintimeout(1),nomaxtimeout,logout(1800),nfstasks(8,8),restimeout(48,0),
hfs(/hfs),bufhigh(32M),readaheadmax(16K),cachewindow(128),percentsteal(20),
maxrdforszleft(32),logicalcache(16M),smf(userfile),pcnfsd,
security(exports,exports,exports),leadswitch,sfmax(0),nochecklist,
fn_delimiter(,),readdirttimeout(30),
public(NFSUSR3,/HFS/usr/lpp/internet/server_root/pub)
```

Ending your z/OS session - mvslogout

Use the **mvslogout** command to disconnect from the remote z/OS NFS server host. The mvslogout command is only required when the **mvslogin** command was used to begin the connection.

An mvslogout to an z/OS user ID cancels a prior mvslogin to the same z/OS user ID from the same local host.

Your account is automatically logged out if it is inactive for the period of time specified in the **logout** site attribute.

The following example disconnects the client from the remote z/OS NFS server machine, *mvshost1*.

```
mvslogout mvshost1
```

Chapter 8. Initialization attributes for the z/OS NFS client

This chapter contains information about the attributes that are used by the z/OS NFS client.

Table 12 contains directive information about this chapter's contents:

Table 12. Attributes - z/OS NFS client

Section	Page
Client attribute syntax	79
Datacaching attribute	84
Mount processing parameters and installation parameters	84

Client attribute syntax

Client attributes are described in Table 13.

Table 13. Client attributes

Attribute	Description
acdirmax(*n***)**	Specifies the maximum lifetime in seconds of cached directory attributes. The **acdirmax** attribute default value is **60**.
acdirmin(*n***)**	Specifies the minimum lifetime in seconds of cached directory attributes. The **acdirmin** attribute default value is **30**.
acregmax(*n***)**	Specifies the maximum lifetime in seconds of cached file attributes. The **acregmax** attribute default value is **60**.
acregmin(*n***)**	Specifies the minimum lifetime in seconds of cached file attributes. The **acregmin** attribute default value is **3**.

© Copyright IBM Corp. 1991, 2006

Table 13. Client attributes (continued)

Attribute	Description
attrcaching(Y | N)	Specifies whether to process attributes and data caching. If attribute caching is in effect, the z/OS NFS client maintains cache consistency with the copy of the file on the NFS server by performing the consistency check with the cached file attributes. When a file's data is read, it remains valid on the z/OS NFS client until the attribute cache is timed out or negated. If **attrcaching(N)** is specified, it will automatically set **datacaching(N)**. The **attrcaching** attribute default value is **Y**.
biod(*n*)	Specifies the number of asynchronous block input/output (I/O) daemons. The BIOD daemon runs on all NFS client systems. When a user on a client wants to read or write to a file on a server, the BIOD daemon sends this request to the server. The BIOD daemon is activated during system startup and runs continuously. The number of daemons is based on the load the client can handle. Six to eight daemons can handle an average load. You must run at least one daemon for NFS to work. The valid range is 1 to 32. The **biod** attribute default value is **6**.
bufhigh(*n*)	Specifies the storage limit for data buffers for the NFS client. The valid range is 1 MB to 128 MB. The **bufhigh** attribute default value is **32** MB.
cln_ccsid(*x*)	Specifies the coded character set identifier (CCSID) for the local mounted file system. The **cln_ccsid** attribute default is **1047** (LATIN OPEN SYSTEM EBCDIC).
convserv(*technique*)	Specifies the conversion technique-search-order which Unicode Service will use for specified **srv_ccsid(*x*)** and **cln_ccsid(*x*)** code pages. *Technique* consists of up to five technique-characters corresponding to the available techniques: R, E, C, L and M. See *z/OS Support for Unicode: Unicode Services* for detailed descriptions on these conversion techniques. The **convserv** default is **LRE**.

Table 13. Client attributes (continued)

Attribute	Description
datacaching(Y | N)	Specifies whether to perform data caching. The **datacaching** attribute provides finer granularity in controlling whether file data should be cached by the z/OS NFS client. By caching the file data, all subsequent references to the cached data is done locally thus avoiding the network overhead. This has additional significance when obtaining data from NFS server systems which do not use UNIX access permissions for security as there is a potential security exposure allowing unauthorized users to access file data. The **datacaching** attribute default value is **Y**.
delaywrite(*n*)	Specifies the maximum number of disk blocks for delay write. The valid range is 0 to 32. The **delaywrite** attribute default value is **16**. The blocksize is 8192. This option is only valid when **datacaching=Y**.
delim (binary | nl | cr | lf | crlf | lfcr)	Specifies the line delimiter for record access to remote files through the basic sequential access method (BSAM), queued sequential access method (QSAM), and virtual storage access method (VSAM). **binary** Specifies the data does not have record delimiters. The access method does not add a delimiter for each record on output and treats any delimiters on input as data. The following text options can be specified: **cr** Specifies that records are delimited by the EBCDIC carriage return character (x'0D'). **crlf** Specifies that records are delimited by the EBCDIC carriage return character followed by the EBCDIC line feed character (x'0D25'). **crnl** Specifies that records are delimited by the EBCDIC carriage return character followed by the EBCDIC new line character (x'0D15'). **lf** Specifies that records are delimited by the EBCDIC line feed character (x'25'). **lfcr** Specifies that records are delimited by the EBCDIC line feed character followed by the EBCDIC carriage return character (x'250D'). **nl** Specifies that records are delimited by the EBCDIC new line character (x'15').
disablella(Y | N)	Specifies the disabling of Lookup Look-Aside (LLA) caching. z/OS NFS server is not compatible with Lookup Look-Aside (LLA) caching enabled. A mount request will fail with client attribute disablella(N) in effect. See *z/OS UNIX System Services File System Interface Reference* for further details on the use of disablella.
dynamicsizeadj(Y | N)	Specifies whether to perform the packet size adjustment for remote procedure call (RPC). The **dynamicsizeadj** attribute default value is **Y**.

Table 13. Client attributes (continued)

Attribute	Description
proto(tcp | udp)	Specifies the transport protocol for the NFS client to communicate with the NFS server. If **secure(udp)** and **proto(tcp)** are both in effect as mount parameters, **proto(tcp)** is ignored. By default, the NFS client selects the **proto** and **vers** with the following priorities: 1. **proto(tcp)** and **vers(3)** 2. **proto(udp)** and **vers(3)** 3. **proto(tcp)** and **vers(2)** 4. **proto(udp)** and **vers(2)**
readahead(*n*)	Specifies the maximum number of disk blocks to read ahead. The block size is 8192 bytes. The valid range is 0 to 16. The **readahead** attribute default value is **1**.
retrans(*n*)	Specifies the number of times to retransmit the NFS remote procedure calls (RPC). The valid range is 0 to 1000. The **retrans** attribute default value is **3**. This option is only valid when **soft** is specified.
retry(*n*)	Specifies the number of times to retry the mount operation. The valid range is 0 to 20,000. The **retry** attribute default value is **10000**. This option is only valid during mount processing.
resize(*n*)	Specifies the read buffer size in *n* bytes. The valid range is 1 to 8192. The **resize** attribute default value is **8192**.
rpcbind(Y | N)	Specifies whether the target NFS server platform supports the RPCBIND protocol, so the NFS client will not have to attempt to use the RPCBIND protocol if that protocol is not supported. The default is **rpcbind(Y)**, to indicate that RPCBIND is supported. If **N** is specified, the z/OS NFS Client will immediately use the PORTMAPPER protocol instead. This keyword has no effect if the client system is not enabled for TCP/IP version 6. The **rpcbind** default value is **Y**.

Table 13. Client attributes (continued)

Attribute	Description
secure(*udp***)**	Specifies the transport protocol for the NFS client to use to bind reserved (privileged) ports when communicating to the NFS server. If **secure(udp)** is specified, **proto(tcp)** is ignored and the NFS client uses udp as the transport protocol to communicate with the NFS server.
soft | hard	Returns an error if the NFS server does not respond or continues to retry the NFS remote procedure call (RPC) until the NFS server responds. • If **hard** is specified, the NFS remote procedure call (RPC) is retried until the NFS server responds. • If **soft** is specified, an error is returned if the NFS server does not respond. The maximum number of retries is specified with the **retrans** option. This option is valid for all NFS RPCs under the mount point.
srv_ccsid(*x***)**	Specifies the coded character set identifier (CCSID) for the remote mounted file system. The **srv_ccsid** attribute default value is **819** (ISO 8859–1 ASCII).
timeo(*n***)**	Sets the remote procedure call (RPC) timeout to *n* tenths of a second. The **timeo** attribute default value is **7**.
vers(2 | 3)	Specifies the NFS protocol version that the client will use to communicate with the NFS server. If no version is specified, the z/OS NFS client will communicate with the NFS server at the highest protocol level that is supported by the server.
wsize(*n***)**	Sets the write buffer size to *n* bytes. The **wsize** attribute default value is **8192**.
xlat(Y | N)	If **Y** is specified, the data in all the files are text and the NFS client will perform data conversion according to the **cln_ccsid** and **srv_ccsid** parameters. The **xlat** attribute default value is **N** and should be used for binary data.

Datacaching attribute

Security checking is done on the Network File System server to determine whether the requesting client user is authorized to access the data. On UNIX systems, this is done by validating the client's user ID and group ID against the file's permission codes. If the authorization checking is successful, the file data is returned to the z/OS NFS client system. Further authorization checking for subsequent access to the cached data or for other client users is done on the z/OS NFS client system.

For conventional MVS data set access through the z/OS NFS server, the user is required to present their z/OS credentials which are checked by the z/OS security system, such as RACF, before file data is returned. Since the z/OS system does not maintain UNIX style permission codes for MVS data sets, the z/OS NFS server returns a code indicating that anyone can access the file. This is done since passing any lesser access code to the client would result in the client user not being allowed to use the cached data which was already read. When the file data is cached on the z/OS NFS client system and another client user on this system attempts to access the same file data, the z/OS NFS client checks the returned permission codes to validate access. Since the z/OS NFS server has passed a code which allows anyone access to the file, all users on the client system can access the cached data without further restrictions. If data caching is turned off, no client caching takes place and each user must pass the server security check.

Based on the installation time out values, the file data cached by the client is flushed and further attempts to access the file data again requires passing server authorization.

The installation **datacaching** parameter can be set and it can be overridden for each mount point so that different mount points can be handled as required for the files under that mount point.

If the potential security exposure can not be tolerated for sensitive file data, the **datacaching** should be used so that no file data is cached by the z/OS NFS client.

Mount processing parameters and installation parameters

Table 14 shows mount processing parameters.

Table 14. Mount processing parameters

acdirmax(n)	**proto(tcp | udp)**
acdirmin(n)	**readahead(n)**
acregmax(n)	**retrans(n)**
acregmin(n)	**retry(n)**
attrcaching(Y | N)	**rpcbind(Y | N)**
cln_ccsid(n)	**rsize(n)**
convserv(UNICODE technique)	**secure(udp)**
datacaching(Y | N)	**srv_ccsid(n)**
delaywrite(n)	**timeo(n)**
delim (binary | nl | cr | lf | crlf | lfcr)	**vers(2 | 3)**
dynamicsizeadj(Y | N)	**wsize(n)**
hard | soft	**xlat(Y | N)**

Table 15 shows installation parameters.

Table 15. Installation parameters

attrcaching(Y | N)	**delim (binary | nl | cr | lf | crlf | lfcr)**
biod(n)	**dynamicsizeadj(Y | N)**
bufhigh(n)	**readahead(n)**
cln_ccsid(n)	**rpcbind(Y | N)srv_ccsid(n)**
convserv(UNICODE technique)	**xlat(Y | N)**
datacaching(Y | N)	
delaywrite(n)	

The following conditions may cause the NFS client to fail its initialization:

- The NFS client is not started in a standalone colony address space.
- The NFS client is already started; multiple instances of the NFS client on a single z/OS system is not supported.
- Invalid parameter is specified in the installation parameters.
- If Unicode exists, then Unicode is used. If Unicode does not exist and Character Data Representation Architecture (CDRA) exists, then CDRA is used. If both Unicode and CDRA do not exist, then initialization fails.

A WTO message is issued to the oeprator console if the NFS client fails to initialize.

NFS client translation support

Table 16 contains NFS client options. See Table 21 on page 101 for more information about considerations for Native ASCII environment support.

Table 16. z/OS NFS clients with non-z/OS based NFS servers

Client Options Specified	Mount Option	Read	Write
xlat(Y), cln_ccsid,srv_ccsid	No TAG specified	NFS client does translation	NFS client does translation
xlat(N)	TAG(TEXT,CCSID)	Logical file system does translation	Logical file system does translation
xlat(Y)	TAG(TEXT,CCSID)	Mount will fail	Mount will fail

Notes:

1. The logical file system will do translation when the **mount tag** option is specified. It will do translation based on the process tag (calling application) and file tag (if the file tag is not zeros or untagged). Otherwise, the system will do translation based on the process tag and **mount tag** for the CCSID information.
2. It is assumed that the user doing the mount knows the files being accessed from the remote non-z/OS file systems. So the CCSID needs to be set accordingly. Data written to the server will be stored in a specific CCSID format. To read it back correctly, the correct CCSID must be specified (for example, without it being translated with the wrong CCSID).

   For more information about client mount options, see Chapter 6, “Commands and examples for AIX and UNIX clients,” on page 47 and Chapter 7, “Commands and examples for z/OS NFS clients,” on page 59.

z/OS NFS client with z/OS NFS server

Both the client and server operate as described in "NFS client translation support" on page 85 and "NFS servers with non-z/OS based NFS clients" on page 101.

In order to avoid double translation, the mount to the server must specify the correct **cln_ccsid** (server option) and the client TSO MOUNT command should not have the **tag** option. The client mount option **xlat(N)** should be specified so that only the server will do translation (if needed) and return the data in the correct CCSID.

```
mount filesystem(NFS001) type(nfs) mountpoint('/u/nfsdir')
          parm('mvsnfs:"/hfs/u/user,text,cln_ccsid(2000)",xlat(N)')
          vi /u/nfsdir/file1

**  Translation will be done by the server only based on file1's file tag and
    cln_ccsid of 2000.
```

In all other cases, double translation may occur as the server will do translation based on its file tag and **cln_ccsid** settings and the logical file system will do translation based on the process tag and the CCSID in the **mount tag** option. Caution must be used as double translation may result in the data becoming garbage.

Chapter 9. Initialization attributes for the z/OS NFS server

This chapter contains information about the attributes that are used to manipulate files in the z/OS NFS server.

Table 17 contains directive information about this chapter's contents.

Table 17. Attributes - z/OS NFS server

Section	Modification	Description	Page
Data set creation attributes syntax	Data set creation attributes can be modified by the client	Data set creation attributes provide information about an MVS file to the z/OS NFS server, such as the type of file, or how the file is allocated (for example, blocks, cylinders, or tracks)	88
Processing attributes syntax	Processing attributes can be modified by the client	Processing attributes provide information to the z/OS NFS server about how to handle the file, such as how long the files remain open, or whether the files are processed in text or binary format	93
Site attributes syntax	Site attributes can only be modified by the system administrator	Site attributes control the z/OS NFS server resources	104

Attributes used for z/OS UNIX System Services file access

These attributes are specific to the following z/OS UNIX file access.

- **hfs***(prefix)*
- **sync** and **async**
- **extlink**

These attributes are relevant to accessing the following z/OS UNIX files as well as conventional MVS data sets.

- **restimeout** - Resource timeout
- **logout** - User log time out
- **security** - Security checking
- **text** - ASCII to EBCDIC data conversion and vice versa
- **binary** - No ASCII and EBCDIC
- **xlat** - Customized translation table

Multipliers

Instead of entering the entire numeric values for the attributes, you can use the multipliers K (1024), M (1024 × 1024), or G (1024 × 1024 × 1024). For example, entering 10M is the same as entering 10,485,760.

© Copyright IBM Corp. 1991, 2006

Duplicate attributes

Specifying an attribute several times on a line does not cause an error. The line is read from left to right, and the last of any duplicate attribute is used. For example:

```
$ vi "file,recfm(vb),recfm(fb)"
```

This results in a file created with a fixed-blocked format.

Data set creation attributes syntax

The data set creation attributes are used to define the structure of MVS data sets when creating a file. These attributes correspond to the data control block (DCB) or the job control language (JCL) parameters used to define an MVS data set when it is created. See *z/OS MVS JCL Reference* for more information about data set creation attributes.

The data set creation attributes do not apply for z/OS UNIX files.

Table 18 describes data set creation attributes. Defaults are underlined **in this format**. You can override these attributes by using the **mount** command or file creation command. For PDS and PDSE, members have the same attributes as the data set attributes, so the file creation attributes specified for members are ignored.

Table 18. Data set creation attributes

Data Set Creation Attribute	Description
blks	Specifies that disk space is allocated by blocks, except for VSAM data sets. See the **space** attribute in this table.
cyls	Specifies that disk space is allocated by cylinders.
recs	Specifies that disk space is allocated by records for VSAM data sets. The **blks** and **recs** attribute values are identical for VSAM data sets.
trks	Specifies that the disk space is allocated by tracks.

blksize(0 | *quan*)
Specifies the maximum length, in bytes, of a physical block on disk. The value of *quan* can range from **0** (the default value) to 32,760. If blksize(0) is specified, the system determines an optimal block size to use.

dataclas(*class_name*)
Specifies the data class associated with the file creation. The *class_name* must be defined to DFSMS/MVS® before it can be used by the client. The system-managed storage automatic class selection (ACS) routine must also assign a storage class to the file being created.

For more information about data classes, see *z/OS DFSMSdfp Storage Administration Reference*.

Table 18. Data set creation attributes (continued)

Data Set Creation Attribute	Description

dir(27 | *quan*)

Specifies the number of 256-byte records needed in the directory of a PDS. Use the **dir** attribute with the **mkdir** command when you create a PDS.

The value of *quan* can range from 1 to 16,777,215. The default value is **27**. The maximum number of PDS members is 14,562.

---

dsntype(library | pds)

Specifies whether a PDSE or a PDS is to be created when the **mkdir** client command is used.

The following options can be specified.

library Specifies partitioned data set extended (PDSE)

pds Specifies partitioned data set (PDS)

You cannot create a PDS (or PDSE) within another PDS (or PDSE).

For more information about creating a PDS or a PDSE, see *z/OS DFSMS Using Data Sets*.

---

dsorg(*org*)

Specifies the organization of a data set.

The following *org* values can be specified.

da Direct access data set

indexed VSAM KSDS data set

nonindexed VSAM ESDS data set

numbered VSAM RRDS data set

ps Physical sequential (ps) data set

The **dsorg** attribute is ignored for directory-oriented client commands.

If you are using VSAM data sets in binary mode, then **nonindexed** is recommended.

---

keys(*len*, *off*)

Specifies the length and offset of the keys for VSAM KSDS data sets. The **keys** attribute can only be specified when using **dsorg(indexed)**.

The *len* and *off* values are specified in bytes.

len Specifies a value between 1 and 255. The default value is **64**.

off Specifies a value between 0 and 32,760. The default value is **0**.

When you create a VSAM KSDS data set, the records you are loading into it must be keyed-sequenced or the write fails. Each write of the data set is treated like a first load, and requires that the records being loaded are in ascending key sequence.

Table 18. Data set creation attributes (continued)

Data Set Creation Attribute	Description
lrecl(8196 | *quan***)** The value of *quan* specifies a record length. 1. Length, in bytes, for fixed-length records. 2. Maximum length, in bytes, for variable-length records. If the **blksize** attribute is specified, the value must be at least 4 bytes less than the **blksize** quantity. The value of *quan* can range from 1 to 32,760. The default value is **8196**.	
mgmtclas(*mgmt_class_name***)** Specifies the management class associated with the file creation. The *mgmt_class_name* must be defined to DFSMS/MVS before it can be used by the client. The system-managed storage automatic class selection (ACS) routine must also assign a storage class to the file being created. For more information about management classes, see *z/OS DFSMSdfp Storage Administration Reference*.	

recfm(*cccc***)**
Specifies the format and characteristics of the records in the data set. The value of *cccc* can be 1 to 4 characters, in one of the following combinations.
`f | fb | fs | fbs`
`u`
`v | vb | vs | vbs`
The following are valid record format characters.
`b` Blocked
`f` Fixed-length records
`s` Spanned for variable records, standard format for fixed records
`u` Undefined-length records
`v` Variable-length records
The **recfm** format characters `v`, `f` and `u` are mutually exclusive. The format character `s` is not allowed for a PDS or PDSE.	

recordsize(*avg***,***max***)** Specifies the average and maximum record size for VSAM data sets. The *avg* and *max* values are specified in bytes. They can each range from 1 to 32,760. The default values are **512** and **4096**, respectively. **These values must be equal for VSAM RRDS.**	

Table 18. Data set creation attributes (continued)

Data Set Creation Attribute	Description
rlse Specifies that unused space should be released from the data set the first time a new data set is closed. For slow clients with long pauses between writes, the **rlse** attribute causes space to be released from the primary extent prematurely. Any additional writes will cause secondary space to be allocated. **norlse** Specifies that unused space should not be released from the data set.	
shareoptions(*xreg***,***xsys***)** Specifies the cross-region and cross-system share options for a VSAM data set. The value of *xreg* ranges from 1 to 4. The value of *xsys* is either 3 or 4. The default values are **1** and **3**, respectively. For more information, see "Sharing VSAM Data Sets" in *z/OS DFSMS Using Data Sets*	
This applies to VSAM data sets only. For spanned records of non-VSAM data sets, see the entry for the **recfm** attribute in this table. **spanned** Specifies that VSAM KSDS or ESDS data sets can contain records that span control intervals (spanned records). **nonspanned** Specifies that data sets do not have spanned records.	
space(*prim[,aux]***)** Specifies the amount of primary and auxiliary space allocated for a new data set on a direct access volume. The value of *prim* is the number (from 0 to 16,777,215) of primary tracks, cylinders, or data blocks in the data set. The value of *aux* (optional) is the number (from 0 to 16,777,215) of additional tracks, cylinders, or blocks allocated if more space is needed. If the **space** attribute is not specified, the default is used. The default values are **100** and **10**, respectively.	
storclas(*class_name***)** Specifies the storage class associated with the file creation. The *class_name* must be defined to the DFSMS/MVS before it can be used by the client. For more information about storage classes, see *z/OS DFSMSdfp Storage Administration Reference*.	

Table 18. Data set creation attributes (continued)

Data Set Creation Attribute	Description
unit(*unit_name***)**	Specifies the unit on which to create a data set. The *unit_name* is a generic or symbolic name of a group of DASD devices. The *unit_name* must be specified as 3390 for extended format data sets. **Notes:** • You cannot create or access tape data sets on an z/OS host using the z/OS NFS server. • You cannot create extended format data sets with the z/OS NFS server, except using ACS routines.
vol(*volser***)**	Specifies the name of the DASD volume to use to store the created data set. The **vol** attribute is the keyword and the value of *volser* represents the volume name. If a data set is system-managed, as determined by the DFSMS automatic class selection (ACS) routines, you can omit this attribute.

Processing attributes syntax

Processing attributes are used to control how files are accessed by the client.

Table 19 describes processing attributes. Defaults are underlined **in this format**. You can override the default processing attributes on the mount command or file processing commands.

Table 19. Processing attributes

Processing Attribute	Description
attrtimeout(*n*)	The time (in seconds) that the data set remains allocated after a **lookup** or **getattr** server operation. The default value of *n* is **120**. The value of *n* can range from 1 to 32,767 (9 hours, 6 minutes, and 7 seconds). **Note:** With NFS version 2 and version 3 protocols, the lookup operation searches for a file in the current directory. If it finds the file, **lookup** returns information on the file's attributes and a file handle pointing to the file. With NFS version 4, neither the file's attributes nor the file handle are returned. The file handle is saved by the server and used as an anchor for accessing the file. The **attrtimeout** value is normally greater than the **readtimeout** or **writetimeout** values.
noattrtimeout	The data set is not deallocated after a **lookup** or **getattr** operation. For more information, see "Timeout attributes" on page 99.
binary	Indicates that the data set is processed between the client and server using binary format and no data conversion occurs between ASCII and EBCDIC formats.
text	Converts the contents in the data set between EBCDIC and ASCII formats. Use this format to share text data between clients and z/OS applications. In **text** mode, the following attributes apply only to conventional MVS data sets: • **blankstrip** and **noblankstrip**. (See the entry for **blankstrip** in this table.) • End-of-line specifiers (**cr**, **crlf**, **lf**, **lfcr**, or **noeol**) are used to indicate the MVS logical record boundary. (See the entry for **lf** in this table. See the **xlat** attribute in this table for customized EBCDIC-ASCII tables.)
blankstrip	With text mode, strips trailing blanks from the end of each record of a fixed-length text file when the file is read. Trailing blanks pad the end of each file or record when a text file is written.
noblankstrip	Does not strip trailing blanks from the end of fixed-length records when a fixed-length text file is read. Does not pad records when writing a text file. The file must be of the correct size or an I/O error is reported to the client. For information about the **text** attribute, see the entry for **binary** in this table. This attribute does not apply to z/OS UNIX files.

Table 19. Processing attributes (continued)

Processing Attribute	Description

With **text** mode, use one of the following end-of-line specifiers.

cr	Carriage Return is the end-of-line terminator.
crlf	Carriage Return followed by Line Feed is the end-of-line terminator (standard DOS).
lf	Line Feed is the end-of-line terminator (standard AIX or UNIX).
lfcr	Line Feed followed by Carriage Return is the end-of-line terminator.
noeol	No end-of-line terminator.

For information about the **text** attribute, see the entry for **binary** in this table.

This attribute does not apply to z/OS UNIX files.

---

cln_ccsid(*n*)
Specifies the coded character set identifier (CCSID) for the remote mounted file system (NFS client) when text is being translated.

The default value of *n* is **819** (ISO 8859-1 ASCII).

The **cln_ccsid** attribute applies only to z/OS UNIX files.

---

executebiton
Sets the execute permission bits in user, group, and other (as reported with the **ls** AIX or UNIX command) for a mount point's files. Use when storing executable or shell scripts on the z/OS system.

This option can only be overridden on a mount point basis — not at a command level.

The **executebiton** attribute does not apply to z/OS UNIX files and can only be used with the **mount** command.

executebitoff
Does not set execute bits in user, group, and other for the mount point's files. This value is normally used in the site file.

---

extlink
Specifies the use of the external link command to create, process, and delete a symbolic link to an MVS data set.

The **extlink** attribute is used with the following commands.

ln -s	Create a symbolic link to an MVS data set.
ls -l	Display the attributes and contents of the symbolic link.
rm	Delete the symbolic link.

The **extlink** attribute only applies to z/OS UNIX file objects.

Table 19. Processing attributes (continued)

Processing Attribute	Description
fastfilesize	Specifies to get the file size from SPF statistics, if it exists, for direct access data sets, PDSs, and non-system-managed data sets.
nofastfilesize	Specifies to read the entire file or member to get the file size for direct access data sets, PDSs, PDSEs, and non-system-managed data sets. Using the **nofastfilesize** attribute might cause a noticeable delay when first accessing very large data sets. These attributes apply to MVS data sets, but do not apply to z/OS UNIX files. For more information, see "Using fastfilesize to avoid read-for-size" on page 314.
fileextmap	Enables file extension mapping. The **fileextmap** attribute can be specified at the file command level for the client platforms that support passing of attributes. See "File name extension mapping" on page 30 for related information.
nofileextmap	Disables file extension mapping.
mapleaddot	Enables mapping of a single leading "." from a client file name to a legal leading "$" on z/OS. The **mapleaddot** attribute should normally be enabled for access by AIX and UNIX clients.
nomapleaddot	Disables mapping of a single leading "." from a client to a leading "$" on z/OS. These attributes do not apply to z/OS UNIX files.
maplower	Enables mapping of lower case file names to upper case when accessing files on z/OS, and back to lower case when sending to the network. This option should normally be enabled for access by AIX or UNIX clients. This option only affects file names (high-level qualifiers and user catalog aliases).
nomaplower	Disables mapping of lower case file names to upper case and back to lower case when using files on z/OS. **Note:** Exports data set entries are not translated to upper case when this option is specified in the default attributes. All **mount** requests are case sensitive. These attributes do not apply to z/OS UNIX files.

Table 19. Processing attributes (continued)

Processing Attribute	Description
mapped	The **mapped** attribute should be specified at the mount or site level when a mixed set of data types is to be processed under a single mount point. The determination of whether the data is to be processed as text or binary depends on the rules that are established in the specified sidefile. If a file extension is not mapped to text or to binary using the sidefiles, then the data will be processed according to what has been specified at the mount or site level (**binary** or **text**). If **binary** or **text** is specified at the file command level, that overrides the **mapped** specification.
mvsmnt	With NFS Version 4, the **mvsmnt** attribute should be specified on user mount commands. In NFS Version 4, mount requests are passed to the server in the form of a PUTROOTFH operation followed by a sequence of **lookup** operations. **mvsmnt** indicates to the z/OS NFS server that the associated **lookup** operation is emulating a mount procedure and causes the z/OS NFS server to write the mount point to the mount handle database (MHDB), so the server can automatically recover the mount point during a server restart.
rdrverf	Enables cookie verifier checking for NFS version 3 and version 4 **readdir** and **readdirplus** requests.
nordrverf	Disables cookie verifier checking for NFS version 3 and version 4 **readdir** and **readdirplus** requests.
readtimeout(*n*)	The **readtimeout** attribute specifies the amount of time in seconds before a data set is released after a read operation. The value of *n* can range from 1 to 32,767 (9 hours, 6 minutes, and 7 seconds). The default value of *n* is **90**. The server closes the file when the file times out. The **readtimeout** attribute does not apply to z/OS UNIX files.
noreadtimeout	Specifies the data set is not deallocated after a read operation. For more information, see "Timeout attributes" on page 99.

Table 19. Processing attributes (continued)

Processing Attribute	Description

The z/OS NFS server uses DFSMShsm™ to recall or delete migrated files. The action that the server takes against migrated files depends on which of the **retrieve** or **noretrieve** attributes is active.

The **retrieve** and **noretrieve** attributes do not apply to z/OS UNIX files.

retrieve
When the **retrieve** attribute is active, the server will recall the migrated file if necessary, upon an NFS_LOOKUP request for the file, depending on the files status.

The server may be able to obtain the migrated files attributes without recall (see "Retrieve attributes" on page 99 for additional information). If not, the recall operation is started by the server. The server waits for the recall operation to complete if the file resides on DASD. If the file does not reside on DASD, the server does not wait for the recall operation to complete and returns a "device not available" message. You can attempt accessing the file again later when the recall is complete.

retrieve(*wait*)
When the **retrieve(wait)** attribute is active, the server waits for the recall to finish.

retrieve(*nowait*)
When the **retrieve(nowait)** attribute is active, the server does not wait for the recall to finish, and immediately returns a "device not available" message. You can attempt accessing the file again later when the recall is complete.

noretrieve
When the **noretrieve** attribute is active, the server does not recall the file, and returns "device not available" upon an NFS_LOOKUP, NFS_READ, or NFS_CREATE request for a migrated data set.

For more information, see "Retrieve attributes" on page 99.

---

setownerroot
Specifies setting the user ID in a file's attributes to root for a superuser.

The **setownerroot** attribute can only be used with the **mount** command and does not apply to z/OS UNIX files.

setownernobody
Specifies setting the user ID in a file's attributes to nobody, for a superuser.

---

sidefile(*dsname***)**
Specifies the name of the data set that contains the rules for file extension mapping purposes. If a sidefile name is specified in the attributes data set, then it is used as the default sidefile for the NFS server. A user can also specify an additional sidefile name during a mount operation to be used along with the default. The mapping rules will first be searched in the sidefile specified during the **mount** command and then the default sidefile is searched. To allow file extension mapping, a sidefile name must be specified either as a default or in the **mount** command. The value of *dsname* is a fully-qualified MVS data set name without quotation marks. The **sidefile** attribute can only be specified at the mount level. See GFSAPMAP in NFSSAMP for sample sidefile and syntax rules.

This attribute does not apply to z/OS UNIX files.

Table 19. Processing attributes (continued)

Processing Attribute	Description
srv_ccsid(*n*)	Specifies the coded character set identifier (CCSID) for the local mounted file system (NFS server) when a new file is being created. The **srv_ccsid** attribute has no effect on the translation of an existing file's data. The default value of *n* (if specified) is **1047** (LATIN OPEN SYSTEM EBCDIC). If the **srv_ccsid** attribute is not specified, new z/OS UNIX files will continue to be created as untagged. The **srv_ccsid** attribute applies only to z/OS UNIX files.
sync	Specifies that data transmitted with the **write** request should be committed to nonvolatile media (for example, DASD) by the server immediately when received.
async	The user can alternatively specify the **async** processing attribute to get improved performance. When this attribute is specified, the data will be committed asynchronously. The **sync** | **async** attribute only applies to z/OS UNIX file objects and only for the NFS version 2 protocol.
tag	Specifies that the newly created files should receive a file tag.
notag	Specifies that the newly created files should not receive a file tag. The tag is set to 0x0000.
writetimeout(*n*)	Specifies the amount of time, in seconds, before a data set is released after a write operation. The value of *n* can range from 1 to 32,767 (9 hours, 6 minutes, and 7 seconds). The default value of *n* is **30**. The server closes the file when the file times out. All cached buffers are forced to disk. Normally **writetimeout** values are kept short because write operations result in exclusive locking. However, for slow client machines with long pauses between writes, you should increase the **writetimeout** value. The **writetimeout** attribute does not apply to z/OS UNIX files.
nowritetimeout	Specifies that the data set is not deallocated after a write operation. For more information, see "Timeout attributes" on page 99.

Table 19. Processing attributes (continued)

Processing Attribute	Description
xlat(*member_name***)**	Specifies how to override the installation default translation table during file processing. The *member_name* is the member name of the PDS or PDSE that contains the customized translation table. The system administrator defines this member name in the attribute data set, and PDS or PDSE in the startup procedure. The **xlat** attribute is ignored if specified on the command line. If a customized translation table is not specified in the attribute file or in the **mount** command, **xlat()** is displayed by the **showattr** client enabling command.

Timeout attributes

The values of the following attributes depend on the settings of the associated site attributes:

attrtimeout, **readtimeout**, and **writetimeout** – These attributes must be within the ranges specified by the **maxtimeout** and **mintimeout** site attributes.

noattrtimeout, **noreadtimeout** and **nowritetimeout** – These attributes are valid only when **nomaxtimeout** is specified in the site attributes.

There are three processing attributes which control when files are timed out: **attrtimeout**, **readtimeout**, and **writetimeout**. The server determines which of these timeouts are in effect based on the last file operation. Thus when an existing file is appended, the file cannot be accessed before it times out in the time specified for **writetimeout** and is released by the server, because write operations result in exclusive locking. Similarly, if a file is read, it is not released before it times out in the time specified for **readtimeout** seconds.

The **readtimeout** and **writetimeout** attributes do not apply to the MVS data set or member being opened by NFS version 4 OPEN operation because there is a stateful CLOSE operation that closes and releases the data set or member.

The **readdirtimeout** site attribute controls the internal **readdir** cache used by directory lookups of MVS conventional data sets to be timed out or discarded based on a customizable value. The default is **30** seconds.

Retrieve attributes

The server deletes the migrated file upon an NFS_REMOVE request for a file, regardless if the **retrieve** or the **noretrieve** attribute is active. Typically, an NFS_REMOVE request is preceded by an NFS_LOOKUP request. If the data set was migrated with DFSMS/MVS 1.2 or below, the retrieve attribute causes a recall because NFS_LOOKUP processing needs to open the data set and read for size. If the data set was migrated under DFSMS/MVS 1.3 and DFSMShsm 1.3, and is SMS managed, its attributes were saved on DASD; therefore it is not always necessary to recall the data set to read for size and the data set may be deleted without recall. If the **noretrieve** attribute is active, the NFS_LOOKUP can return a “device not available” message. If the client code decides to ignore the error and issue the NFS_REMOVE, the migrated file is then deleted.

The UNIX command **ls mvshost** does not issue requests for individual files under the **mvshos**t directory. Migrated files under the **mvshost** directory are displayed, but are not recalled. However, the UNIX command **ls -l mvshost** issues NFS_LOOKUP requests for individual files under the **mvshost** directory.

Mapped keyword processing attribute

Table 20 contains *mapped* and existing keyword information.

Table 20. The mapped keyword and existing keywords

SFMAX	SIDEFILE(NAME)	MAPPED	ACTION
=0	X	Don't care	Data processed using existing rules for binary/text
	I		Server won't come up
	M		MOUNT will fail
1–2000	I + M	SET	File extension used in the MOUNT-specified sidefile and then the site-specified sidefile. If an extension is not found, the existing rules for binary/text will be used.
		NOT SET	Data processed using existing rules for binary/text
	I	SET	File extension used in the site-specified sidefile. If an extension is not found, the existing rules for binary/text will be used.
		NOT SET	Data processed using existing rules for binary/text
	X	Don't care	Data processed using existing rules for binary/text
	M	SET	File extension used in the mount-specified sidefile. If an extension is not found, the existing rules for binary/text will be used.
		NOT SET	Data processed using existing rules for binary/text

Legend:

```
I = sidefile specified in installation table
M = sidefile specified in mount command
X = no sidefile specified
```

Native ASCII processing attributes

The **cln_ccsid(n)** and **srv_ccsid(n)** attributes can be specified either as installation defaults or at mount time for more granularity between different mount points. Unless **srv_ccsid** is specified either as an installation default or at mount time, newly created files will not have any file tag set (i.e., file tag is all zeros). These two attributes affect translation only when text processing is involved and only when an existing file has a nonzero or a nonbinary file tag.

Special attention must be paid to the different server attributes specified. See Table 16 on page 85 and Table 21 on page 101.

Considerations for Native ASCII environment support

For applications running on z/OS V1R2 (and higher), a Native ASCII environment is provided for z/OS UNIX file processing.

In this environment, applications can operate on files in either EBCDIC or ASCII format as well as other data formats defined with a coded character set identifier (CCSID) without translation, provided the data is already defined and stored in the data format wanted.

For the z/OS NFS server to operate properly on z/OS UNIX files in this environment, consider the following important factors:

- Unicode Conversion Services must be installed and set up on the system to let the NFS server use it for text translation.
- Two processing attributes, **cln_ccsid** and **srv_ccsid**, are available for the NFS server for translation purposes as well as for the creation of new files. The **srv_ccsid** attribute determines the CCSID of newly created z/OS UNIX files. If **srv_ccsid** is not specified as an installation default or at mount time, then new files continue to be created as untagged, or with a tag of 0x0000 and the old translation method of using translation tables specified by the **xlat** keyword applies.
- Processing (read/write) of tagged files depends on the different server options specified.

NFS servers with non-z/OS based NFS clients

Table 21 contains NFS server options (file tagging with Unicode Services active).

Table 21. File tagging with Unicode Services active

Server Options Specified	File Tag	Read	Write	Create
text,notag[4]	Untagged or Tag=0x0000 or Tag=0xFFFF	Translation using the current xlat tables	Translation using the current xlat tables	New file created with Tag=0x0000
text,notag	Yes	xlate using Src=FileTag and Tgt=cln_ccsid	xlate using Src=cln_ccsid and Tgt=FileTag	N/A (file exists)
text,cln_ccsid,srv_ccsid,notag[12]	Untagged or Tag=0x0000 or Tag=0xFFFF	xlate using Src=srv_ccsid and Tgt=cln_ccsid	xlate using Src=cln_ccsid and Tgt=srv_ccsid	New file created with Tag=0
text,cln_ccsid,srv_ccsid,notag	Yes[3]	xlate using Src=FileTag and Tgt=cln_ccsid	xlate using Src=cln_ccsid and Tgt=FileTag	N/A (file exists)
text,tag	Untagged or Tag=0x0000 or Tag-0xFFFF	xlate using Src=site attribute srv_ccsid and Tgt=cln_ccsid[10]	xlate using Src cln=ccsid and Tgt=site attribute srv_ccsid[10]	New file created with Tag=srv_ccsid[2]
text,tag	Yes[3]	xlate using Src=FileTag and Tgt=cln_ccsid	xlate using Src=cln_ccsid and Tgt=FileTag	N/A (file exists)
binary,notag[6]	Untagged or Tag=0x0000 or Tag-0xFFFF	No translation	No translation	New file created with Tag=0x0000
binary,notag	Yes[13]	No translation	Fail operation	N/A (file exists)
binary,tag	Untagged or Tag=0x0000 or Tag=0xFFFF	No translation	No translation	If ccsid on mount, Tag=srv_ccsid else Tag=0xFFFF
binary,tag	Yes[13]	No translation	If FileTag!=srv_ccsid fail op else no xlate	N/A (file exists)

Table 21. File tagging with Unicode Services active (continued)

Server Options Specified	File Tag	Read	Write	Create

Notes:

1. Writing to a file that has a tag that is different form the **srv_ccsid** (regardless whether the file is empty or not) will result in the file tag overriding the specified srv_ccsid when text is specified.
2. If srv_ccsid is specified (as an installation default or at mount), then the file is created with the srv_ccsid tag. Otherwise an untagged file is created.
3. xlat is ignored when the file being accessed is tagged.
4. xlat is optional. For untagged files, translation is done using default xlat tables, or custom **xlat** tables (if specified).
5. There is no facility in the NFS server to change an existing file tag. This must be done outside the NFS server.
6. Specifying the binary option overrides any **cln_ccsid** and srv_ccsid specified.
7. All files created by the server when text and srv_ccsid are specified will also have the TXTFLAG set to ON.
8. The NFS file tagging function assumes that Unicode Services is installed and activated on the system (available as of OS/390 V2.8).
9. If Unicode is not enabled, only the NOTAG option is valid
   a. If the TAG option is specified in the site attributes, the NFS server start-up will fail.
   b. If the TAG option is specified on the **mount** command, the mount will fail.
10. File create and write are atomic operations. A file is created before it can be written. Thus, a file always already exists when it is written, and the attributes are used accordingly.
11. For reading or writing untagged files when in text, TAG mode, NFS uses the default server CCSID from the site attribute file. Any **srv_ccsid** values specified on the mount command will be ignored for reading or writing files in this case. The mount srv_ccsid will still be used for file creation however.
12. If TAG is specified in site attribute file, the site attribute **srv_ccsid** and **cln_ccsid** are always used for translating file names.

    If TAG is specified in site attribute file but Unicode Services is not active, the NFS server will shut down.

    If NOTAG is specified in the site attribute file, the site attribute **xlate(table)** is always used for translating file names.
13. These CCSIDs must be specified on the mount.

Table 22 contains NFS server options (file tagging with Unicode Services not active).

Table 22. File tagging with Unicode Services not active

Server Options Specified	File Tag	Read	Write	Create
text,notag	Untagged or Tag=0x0000 or Tag=0xFFFF	Translation using the current xlat tables	Translation using the current xlat tables	New file created with Tag=0x0000
text	Yes	Translation using the current xlat tables	Fail operation	N/A (file exists)
binary,notag	Untagged or Tag=0x0000 or Tag=0xFFFF	No translation	No translation	New file created with Tag=0x0000

Note: The TAG option is not valid if the Unicode Services are not active.

To support the correct use of Unicode Services the new CONVSERV processing attribute is added. The values of this attribute defines the technique search order, or how Unicode Services processes specified code pages.

R - Roundtrip conversion
Roundtrip conversions between two CCSIDs assure that all characters making the "roundtrip" arrive as they were originally.

E - Enforced Subset conversion
Enforced Subset conversions map only those characters from one CCSID to another that have a corresponding character in the second CCSID. All other characters are replaced by a substitution character.

C - Customized conversion
Customized conversions use conversion tables that have been created one on base of special requirements. IBM's National Language Technology Center (NLTC) can provide more information about this tables.

L - Language Environment-Behavior conversion
Language Environment-Behavior conversions use tables that map characters like the *iconv()* function of the Language Environment Runtime library.

M - Modified Language Environment-Behavior conversion
Modified Language Environment-Behavior conversions use tables that map characters like the iconv() function of the Language Environment Runtime library does for converters ending with "C" (for example IBM-932C)..

Default value is "LRE"
This value is recommended to provide correct translation of EBCDIC-newline to ASCII and back.

The value of this attribute should exactly concur with the value of the technique search order, that was used during the current Unicode Image generation.

Site attributes syntax

You can use the site attributes to control z/OS NFS server resources.

Table 23 describes the site attributes. Defaults are underlined **in this format**. Some initial settings are shown, but the system administrator might have changed these settings, so use the **showattr** command to show the actual settings being used. The site attributes cannot be modified by client users.

Table 23. Site attributes

Site Attribute	Description
bufhigh(*n*)	Specifies the maximum storage (in bytes) of allocated buffers before buffer reclamation (see the **percentsteal** attribute in this table) is initiated. The value of *n* is an integer from 1 MB to 2047 MB (the default is **32** MB). If the combined total specified in the **bufhigh** and **logicalcache** attributes is greater than the available storage in the extended private area (implied by the REGION parameter in your procedure) at startup, the server shuts down immediately. A higher number means more caching and potentially better read performance.
cachewindow(*n*)	Specifies the window size used in logical I/O to buffer client block writes received out of order. The value of *n* is a number from 1 to 256 (the default is **112**). The **cachewindow** attribute does not apply to z/OS UNIX files. The suggested value is some small multiple of the number of BIODs running on a client. The general rule in setting the *n* value of **cachewindow(n)** is `n = (( num of BIOD + 1 ) * (client_max_IO_buffer_size/transfer_size))` • *num of BIOD* is the number of blocked I/O daemons set by the client workstation. This value is usually set to defaults at the installation of the operating system or by your system administrator. • *client_max_IO_buffer_size* is the amount of I/O data requested by the client (for example, client writes 8192 bytes of data to the remote file system). This value is determined by your application programs. • *transfer_size* is the actual size of data being sent across the network (for example, the 8192 bytes of data can be broken down to 16 smaller packets of 512 bytes (16x512=8192)). This value is determined dynamically by your client workstation.
checklist	When specified, the server bypasses **saf** checking (even when **saf** or **safexp** is specified) for the list of files and directories underneath mount points which either matches a mount point entry or is a child of a mount point entry in the CHKLIST DD data set. CHECKLIST is only valid if SAF checking is the security option for the particular data access; otherwise, it is ignored even if it is specified. See GFSAPCHK in NFSSAMP library for a sample CHKLIST data set.
nochecklist	When specified, the server operates as before and ignores the information that is specified in the CHKLIST DD data set.

Table 23. Site attributes (continued)

Site Attribute	Description
denyrw	When specified, the server honors deny requests for file share reservations (the Windows Share_Deny value) from the NFS client. The deny requests may be specified on an NFS V4 Open operation or an NLM_share RPC.
nodenyrw	When specified, the server ignores deny requests from NFS clients (the Windows Share_Deny value), and treats the requests as if deny_none were specified.
dhcp	When specified, the server accepts dynamic IP addresses for the NFS client, using the dynamic host configuration protocol (dhcp). The client system must have a static host name and must dynamically update the DNS server with their IP address changes.
nodhcp	When specified, the server supports only NFS clients that use a static IP address.
fileidsize(*n*)	Specifies how to control the handling of fileid sizes by the NFS server in NFS objects. Fileids may be recognized either as 32-bit or 64-bit addresses. Valid values are 32 and 64. The default value is **fileidsize(64)**.

fn_delimiter	Specifies a character 'c' to be used as a delimiter between the file name and the attributes that follow it. This capability allows those sites that have UNIX data sets containing commas to copy and store their data on the NFS server. The following example specifies the default delimiter as a semicolon:
	`fn_delimiter(;)`
	So a user can process a file called 'comma,in-name' by entering:
	`vi "comma,in-name;text,lf"`
	A user can also include a default file name delimiter as a comma as follows:
	`fn_delimiter(,)`

fn_delimiter(,)	The default file name delimiter is a comma.

Table 23. Site attributes (continued)

Site Attribute	Description
hfs(*prefix***)**	Specifies a new z/OS UNIX file system *prefix* to be imbedded in the mount directory path name. The default value of the z/OS UNIX file system prefix is **/hfs**. Mount requests received by the z/OS NFS server beginning with the z/OS UNIX file system prefix value are identified as mount requests for z/OS UNIX. The z/OS UNIX file system prefix value is not part of the path name. **Notes:** 1. The z/OS UNIX file system must already be mounted locally by z/OS UNIX. Otherwise, the client mount request will fail. 2. The prefix value can only be 7 characters or less including the beginning "/"
hfsfbtimeout(*n***)**	Specifies how to control the timeout of the z/OS UNIX vnode token used by the NFS server. The timeout value controls how long before *vnode tokens* saved in file blocks are released. The valid range is 1 to 32,767 seconds. • The value of *n* can go as low as 1 second but to avoid the possibility of the client hanging (because of network delays). The value of *n* is not recommended to be lower than 5 second. • The value of *n* may need to be increased if the network is slow and the accessed directory has a lot of entries. The **hfsfbtimeout** attribute default value is **60** seconds.

Table 23. Site attributes (continued)

Site Attribute	Description

hfssec(*krb5***,***krb5i***,***krb5p***,***sys***)**

Specifies the acceptable network transmission levels of security which can be used as the authentication flavor on NFS version 4 requests for accesses to z/OS UNIX files. This attribute is only used when not overridden by authentication specifications in the exports file. Multiple values for this attribute can be specified using the comma as delimiter. The following are the supported values:

krb5 Provides Kerberos V5 based integrity on the RPC credentials (but not data), when the RPC authentication flavor is RPCSEC_GSS. It uses the DES_MAC_MD5 integrity algorithm and the RPCSEC_GSS service of rpc_gss_svc_none.

krb5i Provides Kerberos V5 based integrity on both the RPC credentials and data, when the RPC authentication flavor is RPCSEC_GSS. It uses the DES_MAC_MD5 integrity algorithm and the RPCSEC_GSS service of rpc_gss_svc_integrity.

krb5p Provides Kerberos V5 based integrity and privacy on both the RPC credentials and data, when the RPC authentication flavor is RPCSEC_GSS. It uses the DES_MAC_MD5 algorithm for integrity and 56 bit DES for privacy. The RPCSEC_GSS service used here is rpc_gss_svc_privacy.

sys Specifies that the AUTH_SYS authentication flavor can also be used to access this file system. Note that the AUTH_SYS authentication flavor does not provide any integrity or privacy protection.

The **hfssec** attribute default is **hfssec(krb5,krb5i,krb5p,sys)**.

Note: file systems that need to have integrity or privacy protection over network transmissions of data, should explicitly specify the desired settings. They should not rely on the default settings because the default settings allow for RPC accesses using the AUTH_SYS authentication flavor which does not provide any integrity or privacy protection.

---

leadswitch

Specifies that the server returns '/' as the first character in each export entry.

noleadswitch

Specifies that the server will not return '/' as the first character in each export entry.

The **leadswitch** attribute is ignored for z/OS UNIX file objects.

---

leasetime(*n***)**

Specifies the length of time (the lease interval) in seconds that the z/OS NFS server allows clients to:

- Reclaim locks and share reservations following an NFS server restart. During this grace period, clients can reclaim locks on behalf of their users.
- Remain active without communicating with the NFS server. If an NFS V4 client does not communicate with the z/OS NFS server for the length of the lease interval, its client id will expire.

The value of *n* can range from 5 to 3600. The specified value must be smaller than the value of the **logout** attribute, if logout is not set to zero. The default value is **120**.

Table 23. Site attributes (continued)

Site Attribute	Description
logicalcache(*n***)**	Specifies the maximum storage (in bytes) for allocated buffers in the logical I/O processing for all the cache windows combined. The value of *n* is an integer from 1 MB to 2047 MB. The default value is **16** MB. The recommended value for this attribute is 50% of the **bufhigh** attribute. The **logicalcache** attribute does not apply to z/OS UNIX files.
logout(*n***)**	Specifies the time limit for inactivity in seconds for a given user on a client. The default value is **1800**. When the limit is reached, the user is automatically logged out. The client user must enter the **mvslogin** command again to reestablish the client's z/OS session. This value should normally be the same as the value defined for TSO/E logout at your site. The value of *n* can range from 61 seconds to 20 megaseconds (approximately 243 days).
maxrdforszleft(*n***)**	Specifies the number of physical block buffers left after determining a file's size. This operation is done for later server read requests to the same file. The buffers left are subject to trimming during a "buffer steal" operation. The value of *n* is an integer from 1 to 1024. The default value is **32**.
maxtimeout(*n***)**	Specifies the maximum timeout allowed. This attribute and the **mintimeout** attribute define the range of values that client users can specify for **attrtimeout**, **readtimeout**, and **writetimeout**. The value of *n* is the number of seconds from 1 to 32,767 (9 hours, 6 minutes, and 7 seconds). This attribute does not affect the **logout** attribute. **nomaxtimeout** Allows client users to specify **noattrtimeout**, **noreadtimeout**, and **nowritetimeout**.

Table 23. Site attributes (continued)

Site Attribute	Description

mintasks(*n*,*m*,*o*)

Defines the minimum number of NFS tasks or threads allowed to run. Tasks may be terminated for reasons such as 878 or 80A ABENDs.

n Specifies the minimum number of subtasks which handle the asynchronous I/O operations or short blocking operations. If the number of active 'short' tasks becomes less than *n* the shutdown process of the NFS server starts.

m Specifies the minimum number of subtasks which handle z/OS UNIX file requests. If the number of active z/OS UNIX tasks becomes less than *m* the shutdown process of the NFS server starts.

o Specifies the minimum number of subtasks which handle long blocking operations. If the number of active legacy long service tasks becomes less than *o* the shutdown process of the NFS server starts.

If *n*, *m*, or *o* are greater than the corresponding values in **nfstasks**, or are not specified, they will default to half the **nfstasks** values.

Valid range for *n* is less than or equal to 99

Valid range for *m* is less than or equal to 100

Valid range for *o* is less than or equal to 99

Valid range for *n* + *o* is less than or equal to 100

---

mintimeout(*n*)

Specifies the minimum timeout. This attribute and **maxtimeout** define the range of values that can be specified for **attrtimeout**, **readtimeout**, and **writetimeout**. The value of *n* is the number of seconds from 1 to 32,767.

The default value is **1**.

---

mixcase

Specifies messages display in mixed case.

upcase

Specifies messages display in uppercase.

Note: Messages GFSA320I, GFSA325I, GFSA349I, GFSA361I, GFSA470I, GFSA471I, GFSA931I, and GFSA934I will always display in uppercase.

Table 23. Site attributes (continued)

Site Attribute	Description
mvssec(*krb5*,*krb5i*,*krb5p*,*sys***)**	Specifies the acceptable network transmission levels of security which can be used as the authentication flavor on NFS version 4 requests for accesses to MVS data sets. This attribute is only used when not overridden by authentication specifications in the exports file. Multiple values for this attribute can be specified using the comma as delimiter. The following are the supported values:. **krb5** Provides Kerberos V5 based integrity on the RPC credentials (but not data), when the RPC authentication flavor is RPCSEC_GSS. It uses the DES_MAC_MD5 integrity algorithm and the RPCSEC_GSS service of rpc_gss_svc_none. **krb5i** Provides Kerberos V5 based integrity on both the RPC credentials and data, when the RPC authentication flavor is RPCSEC_GSS. It uses the DES_MAC_MD5 integrity algorithm and the RPCSEC_GSS service of rpc_gss_svc_integrity. **krb5p** Provides Kerberos V5 based integrity and privacy on both the RPC credentials and data, when the RPC authentication flavor is RPCSEC_GSS. It uses the DES_MAC_MD5 algorithm for integrity and 56 bit DES for privacy. The RPCSEC_GSS service used here is rpc_gss_svc_privacy. **sys** Specifies that the AUTH_SYS authentication flavor can also be used to access this data set. Note that the AUTH_SYS authentication flavor does not provide any integrity or privacy protection. The **mvssec** attribute default is **mvssec(krb5,krb5i,krb5p,sys)**. Note: data sets that need to have integrity or privacy protection over network transmissions of data, should explicitly specify the desired settings. They should not rely on the default settings because the default settings allow for RPC accesses using the AUTH_SYS authentication flavor which does not provide any integrity or privacy protection.

Table 23. Site attributes (continued)

Site Attribute	Description

nfstasks(*n*,*m*,*o*)

Specifies the number of server processes to initiate on startup.

If **nfstasks**(*n*,*m*) is specified, then the following is true.

- The value of *n* is the number of subtasks that handle the asynchronous input/output (I/O) operations or short blocking operations (the maximum number of concurrent NFS server requests).
- The value of *m* is the number of subtasks that handle the long blocking operations (the maximum number of concurrent NFS server recall and z/OS UNIX requests). Increase this value if your server supports lots of active recall or z/OS UNIX clients.

The valid value range for both *n* and *m* is 1 to 100. The sum of *n* and *m* must be less than or equal to 101.

If **nfstasks**(*n*,*m*,*o*) is specified, then the following is true.

- The value of *n* is the number of subtasks that handle the asynchronous input/output (I/O) operations or short blocking operations (the maximum number of concurrent NFS server requests).
- The value of *m* is the number of subtasks that handle z/OS UNIX requests. Increase this value if your server supports lots of active z/OS UNIX requests.
- The value of *o* is the number of subtasks that handle the long blocking operations (the maximum number of concurrent NFS server recall requests). Increase this value if your server supports lots of active recall operations.

The following are valid value ranges for *n*, *m*, and *o*.

- Valid range for *n* is less than or equal to 99
- Valid range for *m* is less than or equal to 100
- Valid range for *o* is less than or equal to 99
- Valid range for *n* + *o* is less than or equal to 100

The **nfstasks** attribute default is **nfstasks(8,16,8)**.

Table 23. Site attributes (continued)

Site Attribute	Description
nlm	Specifies that the initialization of the z/OS NFS server should include starting the NLM and NSM daemons.
nonlm	Specifies that the initialization of the z/OS NFS server should not include starting the NLM and NSM daemons. The system will run without lockd and statd. Specifying **nonlm** does not affect the availability of byte range locking and share reservation support for NFS version 4 protocol access. If **nonlm** is specified, the NLM may not be started after NFS has initialized. If NLM is desired, you must stop and restart NFS after specifying the **nlm** site attribute. The only way to stop NLM is to shut down the NFS server. It is no longer necessary to define the NLM and NSM startup procedures to a z/OS UNIX segment as UID(0) to RACF because the NLM and NSM startup procedures are no longer supported. **Notes:** 1. The lock data sets must always be allocated, even if **nonlm** is specified in the site attributes. 2. The old startup procedures for NLM and NSM are not shipped in z/OS V1R7; these procedures are obsolete and old copies from previous releases should not be used on V1R7 or later releases.
norec878	Specifies that the recovery processing of 878/80A ABENDs will be turned off. That is, if this type of ABEND occurs, the server will shutdown without recovery. It should only be used for debug.
rec878	Specifies that the recovery processing of 878/80A ABENDs will be turned on, and affected tasks will attempt to recover.
pcnfsd	Tells NFS to start the PCNFSD server.
nopcnfsd	Tells NFS not to start the PCNFSD server.
percentsteal(*n*)	Specifies the percent of data buffers which can be reclaimed for use when the **bufhigh**(*n*) limit has been reached. A higher value means a reclaim operation is performed less often, but the cached data is significantly trimmed on each reclaim. This can result in poor read performance because readahead buffers might be stolen. Lower values result in more frequent reclaim operations, but the cached data normal water mark is higher, meaning possibly better performance by reading from cached data. The value of *n* is an integer from 1 to 99. The **percentsteal** attribute default value is **20**.

Table 23. Site attributes (continued)

Site Attribute	Description

public(*legacy_path*,*hfs_path***)**

Specifies the legacy path (MVS conventional data) and HFS path (z/OS UNIX data) that is associated with the public file handle for WebNFS access. The first path, if specified, is the legacy path. The second path is the HFS path and must start with the HFS prefix specified in the **hfs()** keyword.

If the first path is not present, a comma must precede the second path. If the **public** keyword is specified, then one of the paths must be specified. The **public** keyword must be specified after the **hfs()** keyword in the site attribute table. A **lookup** request with the public file handle determines which of the two paths is being referenced by the pathname that is specified. An absolute pathname will tell the server which of the paths is referenced by matching one of the paths specified. A **lookup** request with a relative pathname will be interpreted as a z/OS UNIX request if HFS is active (hfs_path has been provided); otherwise, it is treated as a MVS request.

The **public** attribute default value is **no public path**.

---

pubsec(*krb5*,*krb5i*,*krb5p*,*sys***)**

Specifies the acceptable network transmission levels of security for accesses to public file systems which can be specified as the authentication flavor of the RPC request. This attribute is only used when not overridden by authentication specifications in the exports file. Multiple values for this attribute can be specified using the comma as delimiter. The following are the supported values:.

krb5 Provides Kerberos V5 based integrity on the RPC credentials (but not data), when the RPC authentication flavor is RPCSEC_GSS. It uses the DES_MAC_MD5 integrity algorithm and the RPCSEC_GSS service of rpc_gss_svc_none.

krb5i Provides Kerberos V5 based integrity on both the RPC credentials and data, when the RPC authentication flavor is RPCSEC_GSS. It uses the DES_MAC_MD5 integrity algorithm and the RPCSEC_GSS service of rpc_gss_svc_integrity.

krb5p Provides Kerberos V5 based integrity and privacy on both the RPC credentials and data, when the RPC authentication flavor is RPCSEC_GSS. It uses the DES_MAC_MD5 algorithm for integrity and 56 bit DES for privacy. The RPCSEC_GSS service used here is rpc_gss_svc_privacy.

sys Specifies that the AUTH_SYS authentication flavor can also be used to access file systems. Note that the AUTH_SYS authentication flavor does not provide any integrity or privacy protection.

The **pubsec** attribute default is **pubsec(krb5,krb5i,krb5p,sys)**.

Note: file systems that need to have integrity or privacy protection over network transmissions of data, should explicitly specify the desired settings. They should not rely on the default settings because the default settings allow for RPC accesses using the AUTH_SYS authentication flavor which does not provide any integrity or privacy protection.

Table 23. Site attributes (continued)

Site Attribute	Description
readaheadmax(*n***)**	Specifies the number of bytes to be read to fill internal buffers during read processing to enhance satisfying read requests directly from cache. This reduces the amount of synchronous physical I/O performed for NFS read requests for sequential read file access. It also reduces context switching overhead on NFS read requests by allowing more read requests to be satisfied directly from cache. The value of *n* is an integer from 1 KB to 128 KB (normally 2 to 4 times the common block size used for file access, which is recommended at 8 KB for AIX file activity). The **readaheadmax** attribute default value is **16,384**. Specifying zero (0) will deactivate **readahead**.
readdirtimeout(*n***)**	Specifies the amount of time, in seconds, before the internal readdir cache that is used for MVS conventional data sets is timed out or discarded. The valid range is from 1 to 32,767 (9 hours, 6 minutes, and 7 seconds). The value of *n* can go as low as 1 second, but to avoid the possibility of *client hanging* (due to network delays and staled cache), *n* is not recommended to be lower than 5 seconds. The value of *n* may need to be increased if the network is slow and the accessed directory has a lot of entries. The **readdirtimeout** attribute default value is **30** seconds.
restimeout(*n***,***m***)**	Specifies a retention period and a clock time for the removal of mount points and control blocks that have been inactive longer than the specified retention period. *n* Specifies the resource retention period for mounts and associated resources. If they have been inactive for more than *n* hours, they are removed. The valid range for *n* is 0 to 720 hours (30 days). The default is **48** hours. If *n* is set to 0, the z/OS NFS server does not remove any mount points or associated resources. *m* Specifies the time of day to do the cleanup for mounts and associated resources that have been inactive more than *n* hours. The time of day is specified as a 24 hour local time value. The valid range for *m* is 0 to 23. The default is **0** (that is, midnight). Because cleanup work slows down the server, set *m* so that cleanup work occurs when the server is lightly loaded. If a mount handle is removed by the cleanup activity, the user must do the **umount** and **mount** operations to access the mount point again. The resource cleanup is also done when the server is shutting down.

Table 23. Site attributes (continued)

Site Attribute	Description

security(*mvs[,hfs,public]***)**

Specifies security options for MVS data sets, z/OS UNIX files, and data that is accessed using the public file handle.

mvs Specifies the security option for MVS conventional data sets. The *mvs* parameter is a required parameter.

hfs Specifies the security option for z/OS UNIX files. The *hfs* parameter is an optional parameter.

public Specifies the security option for data that is accessed with the public file handle. The *public* parameter is an optional parameter.

Note: When the optional parameters (*hfs* and *public*) are not specified, they are assigned the same security option as the first parameter.

You can specify the following security options:

exports
Exports list checking. For z/OS UNIX files, checks UNIX permission bits. The UID is obtained from the client RPC request. No SAF checking.

none
Neither SAF checking nor exports list checking. For z/OS UNIX files, checks UNIX permission bits. The UID is obtained from the client RPC request.

saf
SAF checking. No exports checking. For z/OS UNIX files, checks UNIX permission bits. The UID is obtained from the z/OS UNIX segment using **mvslogin**.

safexp
SAF checking and exports list checking. For z/OS UNIX files, checks UNIX permission bits. The UID is obtained from the z/OS UNIX segment using **mvslogin**.

The **security** attribute default is **security(safexp,safexp, safexp)**.

---

sfmax(*n***)**

Specifies the maximum size (in kilobytes) of allocated storage for all of the sidefiles. The value of *n* is an integer from 0 to 2000. The default value is **0** and it signifies that no mapping is allowed on the NFS server. If **sfmax** is set to **0**, specifying the **sidefile** keyword in the attributes data set will cause the server to shut down and specifying the **sidefile** in any subsequent **mount** commands causes the mount to fail because mapping is not allowed on the NFS server. If the amount of storage specified cannot be obtained during server initialization then the server will shut down immediately.

Table 23. Site attributes (continued)

Site Attribute	Description

smf(*level*[,*switch*])	Specifies the level of SMF support and defines whether or not to start SMF record collection at the NFS server startup.
	
	The following *level* options can be specified:
	
	file
	File usage SMF records are to be produced.
	
	none
	No SMF records are to be produced.
	
	user
	User session SMF records are to be produced.
	
	userfile
	Both user session and file usage SMF records are to be produced.
	
	The following *switch* options can be specified:
	
	off
	Activation of SMF records collection can be done manually by issuing the **modify** command. The *switch* parameter is optional.
	
	on
	Activates SMF records collection at the NFS server startup.
	
	The full syntax of the **smf** attribute follows:
	`smf(none|user|file|userfile[,on|off])`
	
	An example of the **smf** attribute follows:
	`smf(user,on)`
	
	An example follows that shows the **smf** attribute when the value of *switch* is **off**:
	`smf(user)`

Part 3. Customization and Operations

Chapter 10. Customization 119
Protecting your programs and files 119
Protecting the server control files 119
Setting up the z/OS NFS authorization 119
Protecting the file system on z/OS with the NFS V4 protocol 120
GSS credential acquisition 121
Security context acceptance 122
Security negotiation 122
Protecting the file system on z/OS with the Security site attribute 123
Unrestricted data access–security(none) . . . 123
Exports list checking–security(exports) . . . 124
SAF checking–security(saf) 124
SAF and exports list checking–security(safexp) 124
SAF checking with checklist processing . . . 125
File system export 125
Authorization of file operations 125
Customizing installation security exits 126
Using UNIX style credentials for authentication 126
Converting data 126
Creating the conversion environment for Unicode Services 127
Configuring NFS servers for multiple TCP/IP stacks 129
Invocation 129
Example procedure to start an NFS server in a multiple server environment 129
User interactions 130
Errors 130
Messages and codes 130
Collecting NFS usage data 130
Configuring the z/OS NFS client 131
Creating the PARMLIB statement for the client 132
Updating z/OS system data sets for the client 132
PARMLIB updates 132
PROCLIB updates 132
Allocating client log data sets 133
Mounting remote file systems 133
Configuring the z/OS NFS server 133
Attributes data set 133
Exports data set 134
When the modified exports data set takes effect 135
Directory statement 135
Checklist data set 138
Mount handle data sets 139
Lock data sets 139
Converting data between ASCII and EBCDIC - NFS V2 and V3 only 140
Customizing the translation table 140
Enabling the xlat processing attribute . . . 141
Updating z/OS system data sets for the server 142
Allocating the z/OS NFS server log data sets 142
Allocating and modifying mapping side file . . 142
Modifying tcpip.ETC.RPC 142
Setting up a user specified port range 143
Configuring a secure z/OS NFS server 144
Using dynamic client IP addressing 147
Installing the client enabling commands 147
Retrieving commands for AIX, Sun Solaris, and Linux 150
Porting the mvslogin, mvslogout, and showattr commands 153
Porting on different compilers and operating systems 154
Compiling 154
Linking 155
Running commands 156

Chapter 11. Network File System operation . . 157
Starting the z/OS NFS client 157
Stopping the z/OS NFS client 158
Starting the z/OS NFS server 158
Starting multiple servers 160
Stopping the z/OS NFS server 160
Starting the z/OS NFS NSM and z/OS NFS NLM 161
Starting component tracing for the z/OS NFS server 161
Entering operands of the modify command for the z/OS NFS server 164
Exportfs operand 166
Freeze operand 166
List operand 166
Mapfile operand 167
Release operand 167
Status operand 168
Unmount operand 169
Unmntall operand 169
Unmnthfs operand 169
Unmntmvs operand 169
Entering operands of the modify command for diagnosis 170
Flushctr operand 170
Flushlog operand 170
Listlock operand 170
Log=msglevel operand 171
Smf operand 172
Version operand 172

Chapter 12. Installation-wide exits for the z/OS NFS server 173
Requirements for NFS 173
Sample link-edit JCL 174
Storage blocks of the server exits 174
Global exit block (GXB) 174
User exit block (UXB) 175
Login installation-wide exit 175
Requirements of the login exit 177
Options of the login exit 177
Structure of the login exit message 177

© Copyright IBM Corp. 1991, 2006

Contents of the login exit parameter list . . . 177
Login exit parameter list. 178
Request codes to the login exit 178
Return codes from the login exit 179
System initialization routine of the login exit 179
Start of new user session routine of the login exit 179
User login request routine of the login exit . . 180
User logout request routine of the login exit . . 180
System termination routine of the login exit . . 181
File security installation-wide exit 181
Requirements of the file security exit 184
Structure of the file security exit message . . . 184
Contents of the file security exit parameter list 184
File security exit parameter list 186
Request codes to the file security exit 186
Return codes from the file security exit. . . . 186
Validate allocate request routine of the file security exit 186
Validate write request routine of the file security exit. 187
Validate read request routine of the file security exit. 187
Return security permissions routine of the file security exit 188

Chapter 10. Customization

This chapter describes how to configure NFS and how to make it available to users. You can perform these tasks to customize NFS:

Table 24. Customizing NFS

Section	Page
Protecting your programs and files	119
Converting data	126
Creating the conversion environment for Unicode Services	127
Configuring NFS servers for multiple TCP/IP stacks	129
Collecting NFS usage data	130
Configuring the z/OS NFS client	131
Configuring the z/OS NFS server	133
Installing the client enabling commands	147

Protecting your programs and files

This section describes security measures that you should take to protect your programs and files in preparation for customizing NFS. These security measures help you protect the server control files, the NFS server and client installations, and the MVS file system. You can customize the NFS security processing and use UNIX style credentials to verify the identity of a client system.

Protecting the server control files

You should protect the following server's control data sets from unauthorized access with RACF, a component of the Security Server for z/OS.

- Attributes file (read by the server at initialization)
- Exports file (read by the server)
- Mount handle data set (read and updated by the server)
- Checklist data set (read by the server)
- Lock data sets (read and updated by the server)
- Kerberos configuration file (read by the server)
- Kerberos keytab (read by the server)

Setting up the z/OS NFS authorization

The following security measures should be addressed when you install the z/OS NFS server and client:

- All programs that come with the z/OS NFS server and client must reside in an APF-authorized program library.
- You need to define the z/OS NFS server and client to resource access control facility (RACF) and assign the necessary level of authority. You do this by defining a RACF user ID for the z/OS NFS server and client. Because the z/OS NFS server and client are run as started tasks, you also need to define an entry in the RACF-started procedures table which associates the z/OS NFS server and client startup procedure names with the previously defined user IDs. For more

© Copyright IBM Corp. 1991, 2006

information about coding and replacing the RACF-started procedure table, see *z/OS Security Server RACF Security Administrator's Guide* and *z/OS Security Server RACF System Programmer's Guide*.

The z/OS NFS server can now be set up with the **trusted** attribute as follows:

```
SETROPTS GENERIC(STARTED)
RDEFINE STARTED mvsnfs.* STDATA(USER(mvsnfs) GROUP(SYS1) TRUSTED(YES))
SETR CLASSACT(STARTED)
SETR RACLIST(STARTED)
ADDUSER mvsnfs OMVS(UID(0))
```

With trusted authority, the NFS server can perform the following tasks:

- Reconstruct the mount points (from the active mount handle data set) upon startup
- Handle mount requests from client prior to user login
- Handle **ls** or **nfsdir** list commands prior to user login
- Be a trusted user during normal operation

During actual remote client file access, the z/OS NFS server first RACROUTEs the remote client's user ID to determine if the remote client is authorized to access the file system. If the remote client is authorized, the z/OS NFS server switches to its own user ID, which has trusted authority, to access the file system.

- You need to define a z/OS UNIX segment for the z/OS NFS client in the RACF user profile as UID(0).

For TCP/IP security information, see *z/OS Communications Server: IP Configuration Guide*.

For z/OS UNIX security information, see *z/OS UNIX System Services Planning*.

Protecting the file system on z/OS with the NFS V4 protocol

The NFS version 4 protocol improves on the NFS version 2 and 3 protocols with stronger authentication and network transmission protection for NFS data. The NFS V4 protections include encryption algorithms for data privacy, multiple protections per file, and the means to negotiate security as NFS clients explore the file system. The NFS V4 protocol provides these protections through the required RPCSEC_GSS security authentication flavor and the SECINFO operation.

The z/OS NFS server enforces these protections, for client RPC requests that use RPCSEC_GSS, through the site attributes **mvssec**, **hfssec**, and **pubsec**. These attributes provide the Kerberos V5 Security Mechanism (RFC1964) subset of the V4 protocol, at the file-system level. The z/OS NFS server also continues to support NFS V2 and V3, as well as V4 protocol requests with the protections provided by the **security** site attribute, as described in "Protecting the file system on z/OS with the Security site attribute" on page 123.

The export data set also contains a security keyword, **sec**, that specifies the Kerberos authentication level that clients must have to access individual files and data sets on the z/OS NFS server. That is, specific export entries can be further constrained with different authentication flavors by using this security keyword. For example, an important export entry can be protected with **krb5p** level set by the security keyword, while other exported entries in the file system can be accessed by all authentication levels which are specified by the **mvssec**, **hfssec** and **pubsec** site attributes. The authentication flavors specified by the **sec** keyword in the export entries should be a subset of the authentication flavors of site attributes **mvssec**, **hfssec** and **pubsec**. In other words, an authentication level is effective only

if it is specified by the site attribute logically AND'ed with the security keyword. For this reason, if the **sec** keyword is not specified, meaning all flavors are on, the authentication level is defaulted to the site attribute **mvssec**, **hfssec** and **pubsec**. For more information on the export security **sec** keyword, see "Exports data set" on page 134.

The **mvssec**, **hfssec**, and **pubsec** attributes let you specify the default network security flavors, and order, that can be used by requests accessing MVS, z/OS UNIX, and public file systems, respectively. These attributes only apply when the RPCSEC_GSS security mechanism is being used for communicating with the NFS client. These site attributes apply to all NFS versions. If you protect a data set with any one of these transmission attributes, NFS V2/V3 requests will get responses of AUTH TOOWEAK unless 'sys' is listed as a valid authentication flavor. NFS V4 requests that do not comply with these protections will get WRONG SEC. For more information on these site attributes, see "Site attributes syntax" on page 104.

Note: Since z/OS NFS only supports RPCSEC_GSS security for NFS version 4, if one of the site attributes is set to require RPCSEC_GSS, then clients using NFS versions 2 and 3, which only support AUTH_SYS security, cannot access those file systems. On the other hand, if the attribute is specified at a mount point, then only that mount point is affected.

To use the NFS V4 RPCSEC_GSS security flavors, the following changes to the security infrastructure are required:

- Kerberos services must be activated on the z/OS system where the NFS server is running. This activation includes the definition of Realms, Inter-Realm relationships, and the Kerberos Principal for the z/OS NFS server. For details on these definitions, see the *z/OS Integrated Security Services Network Authentication Service Administration*.
- The Kerberos principals on NFS clients need to be defined to RACF and assigned a RACF identity. In addition, for Linux clients a principal nfs/hostname.domain should be defined to RACF. This is because Linux Clients use this principal for mounts and some state operations. For further details on defining principals, see *z/OS Security Server RACF Security Administrator's Guide*.
- When acquiring Kerberos tickets, NFS clients must use MD5 Checksum with DES encryption.
- No UNIX style user ID checking will be performed. Clients will always be validated to check for authorization based on their GSS credentials.
- The z/OS NFS server must have READ access to the IRR.RUSERMAP resource in the FACILITY class.

Note: For more information on setting up the z/OS NFS server with RPCSEC_GSS security, see "Configuring a secure z/OS NFS server" on page 144.

GSS credential acquisition

GSS credentials enable the communicating applications to establish security contexts with each other. They can contain multiple cryptographic keys that are required for authentication and message encryption to be performed with different algorithms. The z/OS NFS server uses Kerberos V5 as its security mechanism for acquiring the GSS credentials. The z/OS NFS server initially acquires these credentials during server startup. The z/OS NFS server uses the credentials for accepting the security context requests from NFS clients, and the same credentials may be used for initiating security contexts during RPC callbacks. Therefore, if the Kerberos security server is running on a different machine than the z/OS NFS

server, the Kerberos principal for this z/OS NFS server must be defined in the Kerberos key table identified by the KRB5_KTNAME environment variable.

Note: For more information on setting up the z/OS NFS server with RPCSEC_GSS security, see "Configuring a secure z/OS NFS server" on page 144.

The z/OS NFS server will attempt to acquire the GSS credentials for the maximum credential lifetime but the actual lifetime of credentials will depend on the lifetime of the underlying Ticket Granting Ticket of the Kerberos Security Server, and is not controlled or governed by the z/OS NFS server. On expiration of the server's GSS credentials, client requests will receive the RPCSEC_GSS documented errors and the client is expected to refresh the contexts and retry the requests.

Security context acceptance

A security context is a data structure that contains information about the cryptographic state of a program on the client communicating with the server, and is required for RPC message security services. NFS clients create security contexts with the z/OS NFS server as part of the RPCSEC_GSS protocol of data flow. The z/OS NFS server accepts security context requests subject to the following restrictions and recommendations:

1. The z/OS NFS server does not support channel bindings.
2. The z/OS NFS server never initiates any requests as an agent of NFS clients and therefore recommends that clients do not use credential delegation services while creating security contexts.
3. The z/OS NFS server does not support the Out Of Sequence Detection services of GSS API. It expects NFS clients to have the seq_req_flag parameter turned off on their calls to GSS API gss_init_sec_context.
4. The z/OS NFS server recommends that the clients do not use the Message Replay services of the GSS API. It expects NFS clients to have the replay_det_req_flag turned off on their calls to the GSS API gss_init_sec_context. Note that the z/OS NFS server's implementation of the RPCSEC protocol provides for the protection against replay attacks.
5. The z/OS NFS server does not allow clients to authenticate as anonymous principals.
6. The z/OS NFS server recommends that NFS clients use mutual authentication services during context creation. The z/OS NFS server will still honor context creation requests from NFS clients that are unable to, or choose not to, use mutual authentication services in the GSS-API. However, clients that would require RPC callbacks from the z/OS NFS server have to support accepting security contexts with mutual authentication, because the z/OS NFS server always initiates security contexts with mutual authentication services.

Security negotiation

The NFS version 4 protocol facilitates the use of multiple RPC authentication flavors. The z/OS NFS server supports the Kerberos V5 security mechanism and all the pseudo flavors of the Kerberos security mechanism using the cryptographic algorithms referred to in NFS V4 (RFC3530). To facilitate selection of a particular pseudo flavor, the z/OS NFS server supports security negotiation using the NFS V4 protocol's SECINFO operation. IBM strongly recommends that security negotiation be done by the NFS clients using the SECINFO operation with an RPC authentication flavor of RPCSEC_GSS with the krb5i or krb5p pseudo security flavors.

When responding to SECINFO for security negotiation (when multiple security flavors are present for a file system or file), the z/OS NFS server uses an order of

preference that has RPCSEC_GSS as the most favored flavor followed by AUTH_SYS. For the authentication flavor of RPCSEC_GSS, the z/OS NFS server has krb5, krb5i, and krb5p as its listed pseudo flavors in descending order of preference. NFS clients are, however, free to choose from any one of the z/OS NFS server-supported security flavors for their NFS V4 requests.

NFS V4 clients that decide to use the AUTH_SYS flavor may still have to do an **mvslogin** like their V2/V3 counterparts, depending on the settings of the **security** site attribute.

Note: Because the NFS version 4 protocol does not use the mount procedure, a mount command from an NFS version 4 client will be translated into a PUTROOTFH operation and a series of lookup operations for each component of the mount pathname. A z/OS NFS server cannot accurately ascertain where the mount command will be completed, so both SAF and EXPORT checking are disabled for the z/OS NFS version 4 lookup operation, making name space navigation available without **mvslogin** (as SAF would otherwise require, and EXPORT would otherwise prematurely prevent until an export entry was matched). When non-navigation operations are received by the z/OS NFS server, the proper SAF and EXPORT checking will assure that proper access is given. It is highly recommended that the attribute **mvsmnt** be used with SAF or SAFEXP enabled for a z/OS NFS version 4 server due to the above limitations. Also, if processing attribute **mvsmnt** is not used for a SAF or SAFEXP z/OS NFS version 4 server, security verification issues may arise when a z/OS UNIX file system is being accessed because the z/OS NFS server may use the UNIX permission bits instead of allowing RACF to perform its job as desired.

Protecting the file system on z/OS with the Security site attribute

You can use the **security** site attribute, with the NFS V2, V3, and V4 protocols, to select the level of protection for different types of data access. A different protection level can be specified for MVS data sets, HFS files, and data that is accessed using the public file handle. The attribute used to protect data access is the **security** attribute. The format of the keyword is **security(mvs[,hfs,public])**. The following are the security options: **exports**, **none**, **saf**, and **safexp**. See "Site attributes syntax" on page 104 for syntax rules.) The z/OS NFS server can be configured to handle security in the following ways:

- None
- Exports list checking
- System Authorization Facility (SAF) checking
- Customized installation security exit
- System Authorization Facility (SAF) checking with checklist processing (to bypass SAF for files and directories under selected mount points)
- A combination of these approaches

Note: The UNIX permission checking against the z/OS UNIX hierarchical file system might appear to be inconsistent if the definition of UID and GID is not consistent throughout the domain of the network.

Unrestricted data access–security(none)

If you do not want to restrict data access, you can use the **security(none)** attribute. Neither exports list checking nor SAF checking is done. Client users can access

z/OS files without a z/OS user ID and without using the **mvslogin** command. They simply mount the z/OS file systems that they want to access and unmount when they are finished. For z/OS UNIX files, the UNIX permission bits are checked before access is granted to the client user. The UNIX permission checking is based on the UID from the RPC request.

Note: If UID or GID from the RPC request is zero, it will be mapped to 65534 (nobody) before the UNIX permission checking is performed.

Exports list checking–security(exports)

When you specify **security(exports)** in the attributes data set, the NFS checks the client IP address against the exports list, which is generated from the exports data set, to determine whether or not a mount is to be granted. The NFS also checks the requested directory (or high-level qualifier) to be mounted. For HFS data, it also checks the UNIX permission bits before granting file access to a client user. The UNIX permission checking is based on the UID from the RPC request. Since SAF checking is not done, client users do not need to use the **mvslogin** command.

Note: If UID or GID from the RPC request is zero, it will be mapped to 65534 (nobody) before the UNIX permission is performed.

For more information about the *exports* data set, see "Exports data set" on page 134.

SAF checking–security(saf)

When you specify **security(saf)** in the attributes data set, the Network File System uses RACF or an equivalent product to control access to z/OS file systems. All RACF requests from the server are made through SAF. SAF directs control to RACF, or an equivalent security product, if it is active.

The server uses SAF to validate the z/OS user id and password supplied by the client user. It also uses SAF to validate that the client user is allowed to access z/OS data. A RACF user ID must be defined for each client user that requires access to the server.

For z/OS UNIX data, z/OS UNIX checks the UNIX permission bits before granting file access to a client user. The UNIX permission checking is based on the z/OS UNIX UID obtained by the **mvslogin** process, not the UID passed in by the client on the RPC request. For users accessing z/OS UNIX, their RACF user ID must have an z/OS UNIX segment defined in the RACF profile.

SAF and exports list checking–security(safexp)

When you specify **security(safexp)** in the attributes data set, the NFS checks for both RACF authorization and exports list authorization before granting a client user access to z/OS data. For z/OS UNIX data, z/OS UNIX checks the UNIX permission bits before granting file access to a client user. The UNIX permission checking is based on the z/OS UNIX UID obtained by the **mvslogin** process, not the UID passed in by the client on the RPC request. This is the most restrictive means of limiting DASD access. It requires client users to use the **mvslogin** command.

For more information about the exports data set, see "Exports data set" on page 134.

SAF checking with checklist processing

When you specify **security(saf)** or **security(safexp)** with the **checklist** attribute, the NFS performs SAF as described in "SAF and exports list checking–security(safexp)" on page 124. The only exception to this is that it will not check the files and directories which are underneath the mount points that either match the mount point or the children of the mount points that are in the checklist data set. This allows users that do not have z/OS userids access to data without havings to do a MVSLOGIN while still maintaining SAF checking for data that requires z/OS userids. For more information, see Appendix E, "Sample CHECKLIST data set," on page 329.

File system export

A system administrator issues the **mount** command to an NFS server and makes a remote file system available to the user. The z/OS server keeps a list of file systems and associated access restrictions in an export file. It then compares incoming mount requests to the entries in the file. If a match is found in the export file and the client is authorized for access, then the file system is successfully mounted.

Table 25 shows server processing of a mount request.

Table 25. z/OS server processing of a mount request

Security Option	Export File	z/OS UNIX File	MVS Data Set
none	Not required	No checking exported	No checking exported
saf	Not required	No checking exported	No checking exported
exports	Required	Checking export file	Checking export file
safexp	Required	Checking export file	Checking export file

Note: MVSLOGIN is not required for NFS mount request.

Authorization of file operations

After the file system is mounted, the user performs the normal file operations on the NFS-mounted remote file system. z/OS NFS server adds the z/OS SAF checking in addition to the UNIX file permissions check. For SAF security, the user must perform **mvslogin** to authenticate to z/OS before attempting to do any file access. Otherwise, the user is denied access due to insufficient authorization.

Note: MVS conventional data sets do not support UNIX permission bits in the file attribute structure. By disabling the SAF security, there is no authorization checking for file operation to MVS conventional data set. The UNIX permission bits checking is still performed for z/OS UNIX file operations when the SAF security is disabled.

Table 26 shows server processing of a file request.

Table 26. z/OS server processing of a file request

Security Option	MVSLOGIN	z/OS UNIX File	MVS Data Set
none	Not required	Check file permission bits	No checking
saf	Required***	SAF check***	SAF check***
exports	Not required	Check file permission bits	No checking
safexp	Required***	SAF check***	SAF check***

Table 26. z/OS server processing of a file request (continued)

Security Option	MVSLOGIN	z/OS UNIX File	MVS Data Set
Note: z/OS UNIX segment must be defined for z/OS file operation. (***This does not apply when checklist requirements are satisfied.)			

Customizing installation security exits

You can write installation-wide exits or replaceable modules that customize Network File System security processing, by using product-sensitive programming interfaces provided by the server. Depending on how you code the exit, client users could be required to use the **mvslogin** command even for the **security(none)** and **security(exports)** attributes.

For more information about customizing your installation's security exits see "Login installation-wide exit" on page 175 and "File security installation-wide exit" on page 181.

Using UNIX style credentials for authentication

Authentication is the process of verifying the identity of a client system. This ensures that one client system cannot masquerade as another client system (perhaps with a different set of privileges). Client systems are identified by a set of credentials and authenticated with verification information passed in messages sent to server systems. There are several different conventions for exchanging authentication information in the NFS protocol, including these credentials:

- Null
- UNIX style
- DES-style
- Other, user written

The z/OS NFS server supports the System Authentication flavor of the RPC protocol that employs the UNIX style credentials for all supported NFS protocol versions. For the NFS version 4 protocol, the z/OS NFS server also supports the RPCSEC_GSS authentication flavor, which employs GSS credentials. For its RPCSEC_GSS authentication support, the z/OS NFS server only supports the Kerberos V5 security mechanism. The z/OS NFS client utilizes z/OS UNIX-socket-enabled RPCs to communicate with remote z/OS NFS servers over a TCP/IP network. The credential includes the user ID (UID), group ID (GID), and a list of GIDs the user belongs to. z/OS NFS supports all GID groups specified in the GID group list, which extends support beyond the 16 GID group restriction of the UNIX style AUTH_SYS authentication flavor.

Converting data

The z/OS NFS client supports data conversion defined by the universal character encoding standard known as the Unicode Standard on z/OS V1R2 (and above) when reading data from a remote NFS server or writing data to a remote z/OS NFS server. The Unicode Standard offers character conversion as well as basic case conversion. Within character conversion, characters are converted from one coded character set identifier (CCSID) to another. CCSID information is obtained from the **cln_ccsid** and **srv_ccsid** parameters.

On z/OS releases below V1R2, the z/OS NFS client supports data conversion defined by the Character Data Representation Architecture (CDRA) when reading

data from a remote NFS server or writing data to a remote NFS server. CDRA characters are converted from one CCSID to another. CCSID information is obtained from the **cln_ccsid** and **srv_ccsid** parameters. See *Character Data Representation Architecture Reference and Registry* and *Character Data Representation Architecture Overview* for additional information.

Only single byte to single byte data conversion is supported. For example, if a client file has a CCSID of 437 and a server file has a CCSID of 297, data conversion will occur between USA ASCII format (CCSID 437) and French EBCDIC format (CCSID 297). Single byte to multiple byte conversion (including double byte character set (DBCS)) is not supported and will result in NFS terminating with an error message.

On z/OS releases V1R2 and above, the **cln_ccsid**, **srv_ccsid**, **xlat** , **tag/notag**, and **convserv** parameters identify whether data conversion is performed, and how data conversion is done. These parameters are supported by the z/OS NFS client installation parameter and TSO MOUNT command. The parameters on a TSO MOUNT command override the parameters specified as a z/OS NFS client installation parameter.

Creating the conversion environment for Unicode Services

The z/OS client or server uses Unicode Services to support data conversion on files in either EBCDIC or ASCII formats as well as other data formats that are defined with a CCSID. Once the Unicode Services are installed, a conversion environment has to be set up to create a conversion image that is loaded into storage. The conversion image contains conversion tables that define the data conversion allowed between CCSIDs. See *z/OS Support for Unicode: Unicode Services* for more information.

With the NFS version 4 protocol, text data and metadata are transferred between the server and client in the UTF-8 data format (ASCII text is not transferred directly). z/OS NFS conversion of UTF-8 text data and metadata requires setting up a conversion environment using the z/OS Unicode Services by creating a Unicode conversion image that defines conversion tables with UTF-8 [CCSID 1208].

z/OS NFS requires a minimum set of Unicode conversion images to support data conversion on files in either EBCDIC or ASCII formats. Unicode conversion image generation should include the following CCSID combinations to define EBCDIC, ASCII, and UTF-8 formats:

```
1047,850;  /* EBCDIC -> ASCII */
850,1047;  /* ASCII -> EBCDIC */
1047,819;  /* EBCDIC -> ASCII */
819,1047;  /* ASCII -> EBCDIC */
277,819;  /* EBCDIC -> ASCII */
819,277;  /* ASCII -> EBCDIC */
1208,1047; /* UTF-8 -> EBCDIC */
1047,1208; /* EBCDIC -> UTF-8 */
1208,850; /* UTF-8 -> ASCII */
850,1208; /* ASCII -> UTF-8 */
1208,277; /* UTF-8 -> EBCDIC */
277,1208; /* EBCDIC -> UTF-8 */
1208,819; /* UTF-8 -> ASCII */
819,1208; /* ASCII -> UTF-8 */
```

If Unicode Services is set up correctly, the d uni,all command from system console will show the conversion images with CCSID combinations. An example follows:

```
00- 05.01.28 SYSTEM1 d uni,all
05.01.29 SYSTEM1 CUN3000I 05.01.28 UNI DISPLAY 205
ENVIRONMENT: CREATED 05/14/2001 AT 00.21.17
MODIFIED 05/14/2001 AT 00.21.19
IMAGE CREATED 05/08/2001 AT 01.50.36
SERVICE: CUNMCNV CUNMCASE
STORAGE: ACTIVE 162 PAGES
LIMIT 524287 PAGES
CASECONV: NORMAL
CONVERSION: 01047-00850- 00850-01047-
01047-00819- 00819-01047-
00277-00819- 00819-00277-
00939-00819- 00819-00939-
```

The **convserv** parameter defines how data conversion is performed between CCSIDs by specifying the conversion technique-search-order which Unicode Services will use for specified srv_ccsid(x) and cln_ccsid(x) code pages. Technique consists of up to five technique-characters corresponding to the available techniques (R, E, C, L and M) used to define the technique search order for Unicode Services to process the specified code pages.

The technique-characters, with description, are defined as follows. See *z/OS Support for Unicode: Unicode Services* for more information.

R Roundtrip conversion

Roundtrip conversions between two CCSIDs assure that all characters making the ″roundtrip″ arrive as they were originally.

E Enforced Subset conversion

Enforced Subset conversions map only those characters from one CCSID to another that have a corresponding character in the second CCSID. All other characters are replaced by a substitution character.

C Customized conversion

Customized conversions use conversion tables that have been created to address some special requirements.

L Language Environment-Behavior conversion

Language Environment-Behavior conversions use tables that map characters like the *iconv()* function of the Language Environment Runtime library.

M Modified Language Environment-Behavior conversion

Modified Language Environment-Behavior conversions use tables that map characters like the *iconv()* function of the Language Environment Runtime library does for converters ending with ″C″ (for example IBM-932C).

The **convserve** parameter uses a default value of ″LRE″ which is recommended to provide correct translation of EBCDIC-newline to ASCII and back. The value of this attribute should exactly concur with the value of the technique search order that was used during the current Unicode conversion image generation. See *z/OS Support for Unicode: Unicode Services* for more information on creating a Unicode conversion image.

Configuring NFS servers for multiple TCP/IP stacks

You can perform the following tasks to set up the NFS servers for multiple TCP/IP stacks:

- Invocation
- Example procedure to start a NFS server in a multiple server environment
- User interactions
- Errors
- Messages and codes

This enhancement exploits the ability of the z/OS Communication Server to configure up to eight TCP/IP stacks simultaneously. One NFS server will be allowed to run on each TCP/IP stack. All NFS servers have their own IP-address and work independently of each other with each connecting to a specific transport provider. So each NFS server will use its own unique set of data sets for mount handle database, error log, site attributes, startup procedures, and exports list.

Note: Each NFS server requires its own TCP/IP stack and each NFS server only supports one TCP/IP stack.

The client works with any NFS server as an independent host. At startup, the client selects an NFS server using the servers IP-address or HOST-NAME on the **mount** parameter. On shutdown of one of the NFS servers, all the clients connected to that server will be forced to make new connections with another NFS server and to repeat the startup procedures such as mvslogins, mount connections, etc.

Multiple NFS server support provides an environment on z/OS where applications can have system flexibility by running a NFS server on each LPAR of one z/OS system. This lets you, for example, have a production and test NFS server run on one z/OS system. The use of multiple NFS servers also provides the ability to define separate security-schemes, to separate workload on different NFS servers, and use separate attribute definitions.

Invocation

To run multiple NFS servers on one z/OS system, it is necessary to have a corresponding number of active TCP/IP stacks. The NFS server and TCP/IP startup procedures for each TCP/IP stack should have different names. Each NFS server startup procedure needs to have the following change:

1. Add the **envar** parameter with the _BPXK_SETIBMOPT_TRANSPORT environment variable to point to the TCP/IP startup procedure.
2. SYSTCPD DD statement to point to its TCP/IP stack profile.

See the example in “Example procedure to start an NFS server in a multiple server environment.”

Example procedure to start an NFS server in a multiple server environment

The following contains a sample procedure to start an NFS server in a multiple server environment.

```
//MVSNFS  PROC MODULE=GFSAMAIN,PARMS='INFO',
//           NFSPRFX=MVSNFS,TCPIP=TCPIP.OS390R10
//GFSAMAIN EXEC PGM=&MODULE,
//           REGION=0M,
```

```
//           TIME=1440,
//*     Use environment variable _BPXK_SETIBMOPT_TRANSPORT to set
//*         affinity for a specific TCP/IP stack. TCPIPR10 is name of
//*         start procedure for a selected TCP/IP stack.
//*
//          PARM=('ENVAR("_BPXK_SETIBMOPT_TRANSPORT=TCPIPR10")',
//          '/&PARMS')
//*
//*       Define dataset name of selected TCP/IP stack.
//SYSTCPD  DD   DISP=SHR,DSN=&TCPIP..TCPIP.DATA
//STEPLIB  DD   DISP=SHR,DSN=USER1.LOADLIB1
//SYSPRINT DD   SYSOUT=*
//*
//*SYSMDUMP DD   DISP=SHR,DSN=&NFSPRFX..SYSMDUMP
//OUTPUT   DD   SYSOUT=*
//SYSERR   DD   SYSOUT=*
//*
//* Define dataset names which will be used by current NFS server
//NFSLOG1  DD   DISP=SHR,DSN=&NFSPRFX..DFSLAB11.LOG11
//NFSLOG2  DD   DISP=SHR,DSN=&NFSPRFX..DFSLAB11.LOG21
//NFSATTR  DD   DISP=SHR,DSN=&NFSPRFX..CNTL(NFSATTRV)
//EXPORTS  DD   DISP=SHR,DSN=&NFSPRFX..CNTL(EXPORT1)
//FHDBASE  DD   DISP=SHR,DSN=&NFSPRFX..DFSLAB11.FHDBASE1
//FHDBASE2 DD   DISP=SHR,DSN=&NFSPRFX..DFSLAB11.FHDBASE3
```

User interactions

For more information about configuring multiple TCP/IP stacks, see *z/OS UNIX System Services Planning* and *z/OS Communications Server: IP Configuration Reference*.

Console operators will need to distinguish between different NFS server console messages received from multiple NFS servers on one z/OS system.

System administrators can use the new support to define separate security-schemes, separate workload on different NFS servers, and use separate attribute definitions.

Errors

Existing NFS Server console messages display the start procedure name for a specific NFS server.

Messages and codes

Each NFS Server console message displays the NFS server start procedure as in the following example.

```
GFSA320I(<PROCNAME>) NETWORK FILE SYSTEM SERVER INITIALIZATION FAILED: text
```

Where (<PROCNAME>) is the name of the NFS server start procedure.

Collecting NFS usage data

The z/OS NFS client does not produce any System Management Facilities (SMF) records. However, it does provide the accounting information to z/OS UNIX for SMF recording. z/OS in turn provides the SMF recording services for all physical file systems (PFSs).

The z/OS NFS server does not produce z/OS UNIX SMF records. However, z/OS UNIX provides the SMF recording services for all physical file systems (PFSs).

You can use the SMF records that the z/OS NFS server produces to keep track of how MVS conventional data sets are accessed, and how long each Network File System user session lasts. The z/OS NFS server writes the following SMF records:

Record type-42 subtype 7
This record, written when a file times out, provides the Network File System file usage statistics.

Record type-42 subtype 8
This record, written when a client user logs out of NFS, provides the Network File System user session statistics.

For records containing Internet Protocol (IP) Version 6 addresses, the z/OS NFS server writes a specific type of SMF record. This record type is indicated by a version number of **2** in the smf42psv record field. In these records, the IP address field (smf42cip) is expanded to hold the larger IP V6 address values. For details on these record fields, see Appendix K, "SMF C and assembler header macros," on page 365.

You can control the SMF data collection in the following ways:

- You can use the **smf** site attribute, described in "Site attributes syntax" on page 104, to determine which, if any, SMF statistics are to be collected.
- You can use the **smf=on** or **smf=off** operand of the **modify** command. See "Smf operand" on page 172 for a description of this command, which turns SMF data collection on and off.
- You can generate an SMF report. Use the SMF report sample routine, GFSAPSMF, that can be found in the NFSSAMP library.
- You can use the SMF C and Assembler header macros. See Appendix K, "SMF C and assembler header macros," on page 365 for copies of the C header macro, GFSASSMF, and the Assembler header macro GFSAUSMF. Both header macros contain the mapping for SMF records and can be found in the NFSMAC library.
- Check the SMF setting in the system in SYS1.PARMLIB(SMFPRMnn) for the SMF record type and subtype, where nn is determined by IEASYSmm and the operator command (SET SMF=nn). The Network File System uses the SMF type 42 record, subtypes 7 and 8. You specify **SMF=nn** so the system picks the member of SMFPRM with suffix nn.

You can display the SMF settings with the **d smf,0** operator command.

The SMF write macro, SMFWTM, is used to write the SMF records to the SMF data set. When the server starts, the SMF option is disabled. Therefore, the operator needs to explicitly enable the SMF collection.

For more information about SMF see *z/OS MVS System Management Facilities (SMF)*.

Configuring the z/OS NFS client

This section describes the tasks you can perform to configure the z/OS NFS client. These tasks include creating the PARMLIB statement and updating z/OS system data sets for the client. This section also includes information about allocating client log data sets and mounting remote file systems.

Creating the PARMLIB statement for the client

During z/OS UNIX file system initialization, the z/OS NFS client is started and run in the logical file system (LFS) colony address space. The filesystype parmlib statement for the z/OS NFS client must be present in the SYS1.PARMLIB(BPXPRMxx) parmlib member in order to start the z/OS NFS client. For more information about z/OS UNIX file system reference see *z/OS UNIX System Services File System Interface Reference*.

Updating z/OS system data sets for the client

To accommodate the z/OS NFS client you must update z/OS system data sets PARMLIB, PROCLIB, and the DD statement.

PARMLIB updates

Add the data set defined in the GFSCPROC STEPLIB containing the z/OS NFS client library to the system's APF authorization list (IEAAPFxx). A sample cataloged procedure named GFSCPROC is provided as a member of the sample library NFSSAMP, see "Sample z/OS NFS client startup procedures" on page 356.

Add the filesystype parmlib statement shown in Figure 16 to the z/OS UNIX parmlib member (BPXPRMxx):

```
FILESYSTYPE
   TYPE(NFS)
   ENTRYPOINT(GFSCINIT)
   PARM('installation parms')
   ASNAME(proc_name)
```

Figure 16. Sample filesystype parmlib statement

The name in the TYPE operand must be NFS, otherwise the utility program nfsstat fails.

The operand ENTRYPOINT(GFSCINIT) specifies the entry point for the z/OS NFS client initialization.

The operand PARM('installation parms') specifies the installation parameters for the z/OS NFS client.

The operand ASNAME(proc_name) specifies the procedure name in SYS1.PROCLIB that is used by z/OS UNIX to start the address space in which the z/OS NFS client is initialized.

Note: The proc_name is also used for the name of the address space.

For data integrity and data isolation among different PFSs, the z/OS NFS client is required to start in a separate and standalone colony address space. To start the NFS client in a separate and standalone colony address space, a unique proc_name must be used.

For information about BSAM, QSAM, and VSAM ESDS files, see "BSAM, QSAM, and VSAM ESDS access to remote files" on page 67.

PROCLIB updates

Add the procedure name, *proc_name*, specified in the **ASNAME**(*proc_name*) operand to the system `PROCLIB`.

A sample cataloged procedure named GFSCPROC is provided as a member of the sample library NFSSAMP, see “Sample z/OS NFS client startup procedures” on page 356.

Modify the MVSNFSC procedure and place it in your system PROCLIB. Add the DD statements:

```
NFSCMSG1 as the DD for the primary log data set
NFSCMSG2 as the DD for the secondary log data set
SYSxDUMP as the DD for the SYSxDUMP data set ('x' = U or M)
```

Allocating client log data sets

For information about allocating the z/OS NFS client primary and secondary log data sets, see Appendix L, “Capturing diagnostic information using z/OS NFS log data sets and from other components,” on page 379.

Mounting remote file systems

z/OS UNIX does not support z/OS NFS mounts in the SYS1.PARMLIB member statement. You can use the z/OS UNIX automount facility (/etc/rc shell scripts support) or the TSO MOUNT command to make a connection between a mount point on your local z/OS UNIX file system and one or more files on a remote MVS, AIX, UNIX, z/OS, or other file system. The remote file system can be mounted using the TSO MOUNT command only after the z/OS UNIX file system initialization is complete. The TSO MOUNT command can only be used by a z/OS UNIX superuser. For more information about the TSO MOUNT command, when used with the z/OS NFS client, see “Mount command syntax and examples” on page 65.

When using the automount facility of z/OS UNIX System Services, the remote file system is mounted on its first data access attempt if it is not already mounted.

When the automount facility is used to manage remote NFS mount points, the z/OS NFS user could experience ESTALE/EIO errors if the automounter unmounts the accessed mount point when the time limits specified by the automount duration and delay parameters have been exceeded. The automount default, DURATION=NOLIMIT, disables the DURATION timeout period. The DELAY for unmounting file systems should be longer than the time z/OS NFS clients are likely to keep z/OS NFS mounts to the files and directories active. For more information about the z/OS UNIX automount facility (/etc/rc shell scripts support) see *z/OS UNIX System Services Planning* and *z/OS UNIX System Services File System Interface Reference*.

The remote file system must be mounted on the z/OS UNIX file system prior to any reference being made to the remote data. Once mounted, the remote file system can be treated as an extension of the local z/OS UNIX file system.

Configuring the z/OS NFS server

This section describes the tasks you can perform to configure the z/OS NFS server. These tasks include allocating and modifying data sets and mapping side files, allocating mount handle, lock, and z/OS NFS server log data sets, modifying tcpip.ETC.RPC, and updating z/OS system data sets for the server. This section also includes information about the data conversion between EBCDIC and ASCII.

Attributes data set

To allocate and modify the attributes data set, perform the following tasks:

1. Allocate a fixed-block partitioned data set or a fixed-block sequential data set with a record length of 80.
2. Copy the sample member GFSAPATT from the *prefix*.NFSSAMP data set into the allocated attributes data set.
3. Modify the attributes to suit your installation. Appendix F, "NFS system server sample attribute table," on page 331 shows the sample attributes data set. You can specify three sets of attributes within the attributes data set.
   - Data set creation attributes
   - Processing attributes
   - Site attributes.

Notes:

1. Client users can override the processing and data set creation attributes (for their own sessions), but not the site attributes.
2. The attributes data set specified on the NFSATTR DD statement of the mvsnfs startup procedure is read during server startup processing. Further changes to this data set do not take effect until the server is restarted. Also, whenever any attributes are changed, all the previous existing mount points have to be unmounted and mounted again (using the **umount** and **mount** client commands) if you want the mount points to pick up the new attributes.

Specify this attributes data set for the NFSATTR DD statement in the MVSNFS cataloged procedure.

Coding attribute statements

Here are guidelines for coding attribute statements:

- You can continue a line in the attributes data set by placing a "\" or "+" at the end of the line.
- A "#;" anywhere in the data set indicates a comment that extends to the end of the line (unless the **altsym** keyword is used in the **start** command or the server startup procedure; if **altsym** is used, a ";" indicates a comment).
- If you specify more than one attribute on a line, separate the attributes with a comma and a space.

Exports data set

To allocate and modify the exports data set, perform the following tasks.

1. Allocate a fixed block partitioned data set or a fixed block sequential data set with record length of 80.
2. Copy the sample member GFSAPEXP from the *prefix*.NFSSAMP data set into the allocated data set.
3. Modify the sample exports data set to suit your installation. Appendix G, "Sample exports data set," on page 349 shows the sample exports data set.
4. Specify this exports data set for the EXPORTS DD statement in the MVSNFS cataloged procedure.

The exports data set contains entries for directories that can be exported to clients. It is used by the server to determine which data sets' high-level qualifiers or z/OS UNIX directories can be mounted by a client in a read or write mode.

Note: You cannot specify exporting a "parent directory" or a subdirectory of an exported directory. For example, if you specify DIR1 in the exports data set, DIR1 and all its subdirectories are exported. You cannot specify any subdirectories under DIR1 in the exports data set.

When the modified exports data set takes effect

The exports data set specified in the exports DD statement of the MVSNFS startup procedure is read during server startup processing. Changes to this file do not take effect for new mount points until the **exportfs** operand of the **modify** command is completed or the server is restarted. The changes will also be immediately enforced at z/OS NFS server startup or EXPORTFS operator command time regardless of the state of any pre-existing mount points.

When the exportfs operand of the modify command is issued, any errors found in the file are sent to the system log, the entire exports list is not refreshed, and processing continues.

When the server is started, if any errors are found in the file, they are sent to the system log, and the server stops once the entire exports data set has been read.

Directory statement

Use the directory statements in the exports data set to limit access of directories to specified client workstations.

- Entries can be up to 4096 characters long. Directory names for MVS data sets must follow z/OS naming conventions.
- Lines can be continued by placing a backslash (\) or a plus sign (+) at the end of the line.
- A "#" anywhere in the data set indicates a comment that extends to the end of the line (unless the **altsym** keyword is used in the **start** command or the server startup procedure; if the altsym keyword is used, the ";" indicates a comment).
- Spaces cannot be used in the keywords.
- The parameters to the right of directory are optional. Except for **ro** and **rw**, the parameters can be combined. If they are combined, only the first parameter is preceded by "-". For example:

```
user1.test -access=rs60001:sun1:sun2,ro
```

An entry for a directory is specified in Figure 17 on page 136.

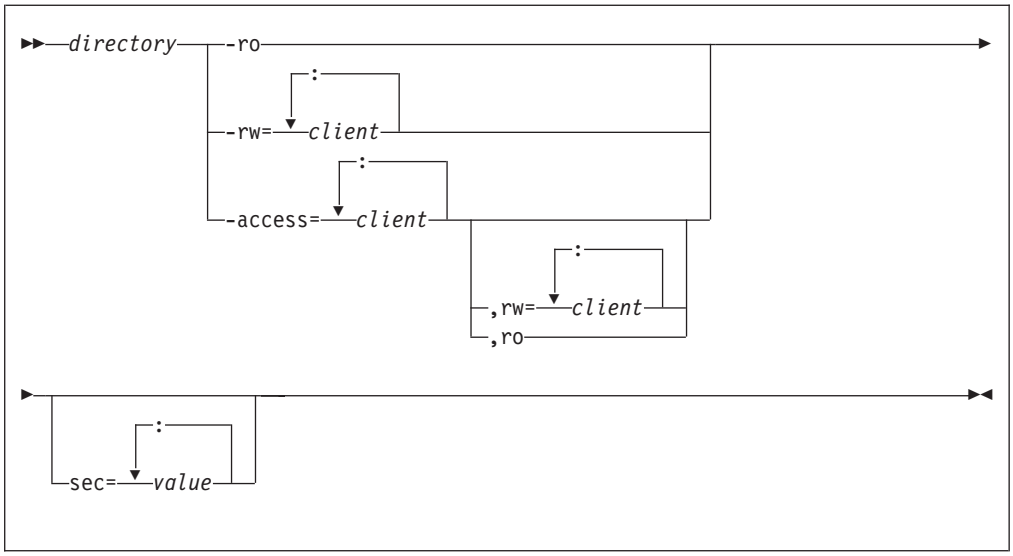

Figure 17. Entry for a directory

where

directory
MVS high-level qualifier, data set name, or alias to a user catalog; the name must conform to MVS data set naming conventions unless z/OS UNIX is used (for a z/OS UNIX directory entry, you need to use the hfs prefix that is used at your site).

-ro
Export the directory as read only. If not specified, the directory is exported as read/write.

-rw=*client***[:***client***...]**
The directory is exported as read/write to specified *clients*, and read-only to everyone else. Use a colon (:) to separate client names.

-access=*client***[:***client***....] [[,rw=***client***[:***client***...]] | [,ro]]**
Gives access only to *clients* listed.

If neither **rw** nor **ro** is specified for the **-access** parameter, then the clients listed have read/write access and the rest of the clients have no access.

If the **rw** parameter is specified for the **-access** parameter, the associated clients have read/write access to the directory, and the clients specified in the **access** list but not in the **rw** list has read-only access.

If the **ro** parameter is specified for the **-access** parameter, the clients in the **access** list have read-only access to the directory, and the rest of the clients have no access.

Use a colon (:) to separate client names.

sec=krb5:krb5i:krb5p:sys
Specifies the acceptable level of network transmission security access that a client's RPC request must provide. If a client attempts to access an object in the directory using a network security level that is not specified on the **sec** parameter, the access is denied.

krb5
Provides Kerberos V5 based integrity on the RPC credentials (but not

data), when the RPC authentication flavor is RPCSEC_GSS. It uses the DES_MAC_MD5 integrity algorithm and the RPCSEC_GSS service of rpc_gss_svc_none.

krb5i
Provides Kerberos V5 based integrity on both the RPC credentials and data, when the RPC authentication flavor is RPCSEC_GSS. It uses the DES_MAC_MD5 integrity algorithm and the RPCSEC_GSS service of rpc_gss_svc_integrity.

krb5p
Provides Kerberos V5 based integrity and privacy on both the RPC credentials and data, when the RPC authentication flavor is RPCSEC_GSS. It uses the DES_MAC_MD5 algorithm for integrity and 56 bit DES for privacy. The RPCSEC_GSS service used here is rpc_gss_svc_privacy.

sys
Specifies that the AUTH_SYS authentication flavor can also be used to access this file system. Note that the AUTH_SYS authentication flavor does not provide any integrity or privacy protection.

Use a colon (:) to separate security levels.

If the **sec** parameter is provided, it further restricts the network transmission protection of specific export entries, in the domain of the file system wide network transmission protection, which is controlled by the **mvssec**, **hfssec**, or **pubsec** site attributes. When this parameter is not specified, the allowed security levels are governed by the **mvssec**, **hfssec**, and **pubsec** site attributes.

Note: If no options are specified, the default value allows any client to mount the given directory with read/write access.

Following are examples of entries in an *exports* data set.

```
mvsnfs          -ro                            # give read-only access
                                               # to all clients
                                               #
theresa.text                                   # give read/write
                                               # access to all clients
                                               #
robert.mixds   -rw=fsrs001:fslab004:fslab007 #
                                               # give read/write access
                                               # to the clients named
                                               # fsrs001, fslab004 and
                                               # fslab007, and give
                                               # read-only access to
                                               # all other clients
                                               #
/hfs/newproductdirectory  -rw=johnson          # give read/write access to
                                               # this z/OS UNIX directory to the
                                               # client named johnson;
                                               # give read-only access to
                                               # all others
                                               #
barbara.pds    -access=fsrs001:fslab007        #
                                               # give read/write access
                                               # only to clients named
                                               # fsrs001and fslab007
                                               #
daniel.pds2    -access=fslab004,ro             # give read-only access
                                               # only to the client
                                               # named fslab004
                                               #
```

```
virginia.vsam    -access=fslab004:fslab007,rw=fslab004 #
                                                 # give read-only
                                                 # access only to the
                                                 # client named fslab007,
                                                 # and give read/write
                                                 # access to fslab004.
                                                 #
/hfs/u     -sec=krb5:krb5i:krb5p               # client must use krb5, krb5i or
                                             # krb5p authentication levels to
                                          # access server, provided that
                                          # hfssec also allows these
                                          # authentication levels.
```

Notes:

1. If your installation cannot use the "#" as a comment delimiter, see "Starting the z/OS NFS server" on page 158.
2. The keywords **ro** and **rw** are mutually exclusive.
3. The ability to write (that is, **rw** specified or **access** specified without other parameters) implies read access also.
4. If **access** and **rw** are specified together, the client names in the **rw** list are logically or'ed with the **access** list to determine the total list of clients with read access.
5. Multiple lines can be used in the exports data set for a given directory to merge the **access** list and the **rw** list. However, similar clauses (for example, an access followed by an access) completely replace any previous specification. If **ro** is specified for a data set on one line and a further line specifies **rw** for that data set, the **rw** undoes the **ro** specified earlier. Similarly, a line with null options completely undoes all previous specifications for that directory, giving read/write access to all clients.
6. It is not appropriate to have the same data set or z/OS UNIX directory defined more than once in the exports data set. If for any reason this is the case, only the last definition in the exports data set is valid.

Checklist data set

To allocate and modify the checklist data set, perform the following tasks:

1. Allocate a fixed block partitioned data set or a fixed block sequential data set with record length of 80.
2. Copy the sample member GFSAPCHK from the *prefix*.NFSSAMP data set into the allocated data set.
3. Modify the sample checklist data set to suit your data security requirements. Appendix E, "Sample CHECKLIST data set," on page 329 shows the sample checklist data set.
4. Specify this checklist data set for the CHKLIST DD statement in the MVSNFS catalogued procedure.

The checklist data set contains entries for prefixes (directories) or file names which should match subsequent mount points or be the parent of the subsequent mount point to allow SAF checking to be bypassed for everything underneath that mount point.

The checklist data set specified in the CHKLIST DD statement of the MVSNFS startup procedure is read during server startup processing. Changes to this file do not take effect for new and active mount points until the server is restarted.

Note: Proper care must be used to protect this data set since it can be used to bypass SAF checking for selected files and directories.

Mount handle data sets

The mount handle data sets are used to record active mounts during server operation and allow clients to stay mounted when the server is shut down and restarted. The Network File System alternates between two data sets to record this information; only one data set is used at a time, and it is switched at either shutdown or at resource cleanup timeout.

To create the mount handle data sets, perform the following tasks:

1. Allocate two empty VSAM KSDS data sets with the following attributes:
   - Starting with offset 0, the first 16 bytes in the record are the prime key field.
   - The maximum record length of the mount handle data set is 2000 bytes, and the average record length is 1700 bytes.

     ```
     DEFINE CLUSTER (NAME(mount_handle_data_set_name) -
                VOL(vsam_volume_name) -
                CYL(1 1) -
                INDEXED -
                REUSE -
                KEYS(16 0) -
                SHAREOPTIONS(1 3) -
                RECSZ(1700 2000))
     ```

     See Appendix M, "GFSAMHDJ sample code for creating NFS mount handle data sets and lock data sets," on page 387 for sample JCL showing how to create the mount handle data sets.
2. Specify these two data sets to the FHDBASE DD statement and FHDBASE2 DD statement in the MVSNFS procedure (see Figure 18).
3. The server switches data sets after resource cleanup has run.
4. Resource cleanup is done at Network File System shutdown and resource cleanup timeout.

Figure 18 shows how to specify the mount handle data set on the FHDBASE DD statement and FHDBASE2 DD statement on the MVSNFS procedure.

```
//*     FHDBASE AND FHDBASE2 ARE
//*     THE MOUNT HANDLE DATA SETS.
//*     THEY NEED TO BE PREALLOCATED
//*     AS EMPTY VSAM KSDS DATA SETS.
//*     THEY WILL BE USED ALTERNATELY.
//*     SAMPLE JCL CAN BE FOUND IN HLQ.NFSSAMP(GFSAMHDJ).
//*
//FHDBASE  DD   DISP=SHR,DSN=MVSNFS.FHDBASE
//FHDBASE2 DD   DISP=SHR,DSN=MVSNFS.FHDBASE2
```

Figure 18. Specifying the mount handle data set in the MVSNFS procedure

Lock data sets

Before z/OS V1R7, lock data sets for the z/OS Network File System were HFS files that contained only client host names, which provided no coherency checking for locks after a server restart. To solve this problem, in z/OS V1R7 the lock data sets are VSAM key-sequenced data sets that record the client host IP address and, for NFS V4, the client identification as well as the client host name. Following a server outage, the z/OS NFS server reads the new lock data sets during initialization to determine which clients can reclaim locks. The z/OS Network File System alternates between two data sets to record this information; only one data set is used at a time, and it is switched at shutdown, at resource cleanup timeout, and at the end of the grace period at server startup.

Note: The lock data sets must always be allocated, even if *nonlm* is specified in the site attributes.

To create the lock data sets, perform the following tasks:

1. Allocate two empty VSAM KSDS data sets, on separate DASD volumes to reduce the possibility that a single failure would result in the loss of both data sets. Allocate them with the following attributes:
   - Starting with offset 0, the first eight bytes in the record are the prime key field.
   - The maximum record length of the lock data set is 2000 bytes, and the average record length is 1700 bytes.

     ```
     DEFINE CLUSTER (NAME(lock_data_set_name) -
                VOL(vsam_volume_name) -
                CYL(1 1) -
                INDEXED -
                REUSE -
                KEYS(8 0) -
                SHAREOPTIONS(1 3) -
                RECSZ(1700 2000))
     ```

     See Appendix M, "GFSAMHDJ sample code for creating NFS mount handle data sets and lock data sets," on page 387 for sample JCL showing how to create the lock data sets.
2. Specify these two data sets to the LDBASE DD statement and LDBASE2 DD statement in the MVSNFS procedure (see Figure 19).
3. The server switches data sets after resource cleanup has run.
4. Resource cleanup is done at Network File System shutdown, resource cleanup timeout, and at the end of the grace period at server startup.

Figure 19 shows how to specify the lock data set on the LDBASE DD statement and LDBASE2 DD statement on the MVSNFS procedure.

```
//*     LDBASE AND LDBASE2 ARE
//*     THE LOCK DATA SETS.
//*     THEY NEED TO BE PREALLOCATED
//*     AS EMPTY VSAM KSDS DATA SETS.
//*     THEY WILL BE USED ALTERNATELY.
//*     SAMPLE JCL CAN BE FOUND IN HLQ.NFSSAMP(GFSAMHDJ).
//*
//LDBASE  DD   DISP=SHR,DSN=MVSNFS.LDBASE
//LDBASE2 DD   DISP=SHR,DSN=MVSNFS.LDBASE2
```

Figure 19. Specifying the lock data set in the MVSNFS procedure

Converting data between ASCII and EBCDIC - NFS V2 and V3 only

For text processing mode, with the NFS version 2 and 3 protocols, data is converted between EBCDIC and ASCII. No double byte character set (DBCS) or multiple byte character set (MBCS) forms of data are converted.

Customizing the translation table

You can customize the translation table for the Network File System using the processing attribute **xlat(member_name)**. The parameter **(member_name)** is the member name of a PDS or PDSE containing the customized translation table. This attribute can be specified either in the mount operation or in the attribute file. It can be specified on a file operation but is ignored, only the mount or the xlat value takes effect.

If the processing attribute, **xlat**, is not specified in the attribute file, the Network File System internal translation table is used as the installation default translation table. When the **xlat(member_name)** processing attribute is specified in the attribute file, this customized translation table becomes the installation default translation table. The Network File System internal translation table is derived from EBCDIC code page 0037 and ISO 8859-(ASCII). RPC arguments, such as filenames, are always translated by the installation default translation table. Data shipped with RPCs are translated by the mount specified translation table, if any. Otherwise, they are also translated by the installation default translation table.

A mount request with processing attribute, **xlat**, specified overrides the installation default translation table.

When accessing z/OS UNIX files, you must specify the OEMVS311 translation table or your customized translation table either in the mount request or in the default attributes. The OEMVS311 table translates ASCII (ISO 8859-1) to and from EBCDIC (1047 - z/OS UNIX). TCP/IP for MVS version 3.1 provides the OEMVS311 table. This table translates the UNIX line terminator (**lf**) to the z/OS UNIX line terminator (**nl**).

See *z/OS Communications Server: IP Configuration Reference* for more information about creating and customizing your own translation tables.

Enabling the xlat processing attribute

A DD statement, NFSXLAT, is required in the Network File System startup procedure to enable the **xlat(member_name)** processing attribute:

```
//NFSXLAT         DD        DSN=dataset_name,DSP=SHR
```

where

dataset_name
: Specifies the name of a PDS or PDSE whose member contains the customized translation table.

A PDS or PDSE, dataset_name, is created by the CONXLAT utility whose member contains the customized translation table.

Notes:

1. See *z/OS Communications Server: IP Configuration Reference*, "Using Translation Tables," for more information about creating and customizing your own translation tables.
2. You can edit or modify the translation table from your own or from a member in the tcpip.SEZATCPX data set and then use CONVXLAT utility to convert the source table into binary format. The CONVXLAT utility can take a PDS or PDSE as input, and its output data set can be physical sequential, PDS or PDSE.
3. The Network File System only supports PDS and PDSE. A sequential data set must be copied to either a PDS or PDSE member.
4. The Network File System does not support the translation for multiple-byte character sets.
5. Sample steps for creating the xlat member:
   a. Run the TCPIP CONVXLAT utility to create a physical sequential (PS) data set with DSORG=PS, RECFM=F, LRECL=256, BLKSIZE=256;

      ```
      "convxlat" 'tcpip.sezactcpx(standard)' 'hlq.xlat.output'
      ```

b. Allocate a PDS data set with DSORG=PO, RECFM=F, LRECL=256, BLKSIZE=256; copy the CONVXLAT output data set as a member in the PDS data set

c. Allocate the xlat member in the z/OS NFS startup procedure.

Updating z/OS system data sets for the server

Update the following z/OS system data sets to accommodate the z/OS NFS server:

- PARMLIB updates

  Add the data set defined in the GFSAPROC STEPLIB containing the z/OS NFS server library to the system's APF authorization list (IEAAPFxx). A sample cataloged procedure named GFSAPROC is provided as a member of the sample library NFSSAMP, see "Sample z/OS NFS server startup procedures" on page 353.

- PROCLIB updates

  A sample cataloged procedure named GFSAPROC is provided as a member of the sample library NFSSAMP, see "Sample z/OS NFS server startup procedures" on page 353.Modify the MVSNFS procedure and place it in your system PROCLIB. Add the DD statements:

```
EXPORTS as the DD for the exports data set
NFSATTR as the DD for the attributes data set
FHDBASE and FHDBASE2 as the DD for the mount handle data set
NFSXLAT as the DD to enable the xlat processing attribute
NFSLOG1 as the DD for the primary log data set
NFSLOG2 as the DD for the secondary log data set
SYSxDUMP as the DD for the SYSxDUMP data set ('x' = U or M)
CHKLIST as the DD for the checklist data set (if needed)
LDBASE and LDBASE2 as the DD for the lock data sets
```

Allocating the z/OS NFS server log data sets

For information about allocating the z/OS NFS server primary and secondary log data sets, see Appendix L, "Capturing diagnostic information using z/OS NFS log data sets and from other components," on page 379.

Allocating and modifying mapping side file

For information about allocating and modifying mapping side files, see Appendix F, "NFS system server sample attribute table," on page 331.

Modifying tcpip.ETC.RPC

Add the entries in Table 27 to the tcpip.ETC.RPC file for the services provided by the z/OS NFS server:

Table 27. Modifying tcpip.ETC.RPC

Service	Number	Description
nfsd	100003	Network File System daemon
mountd	100005	Mount daemon
mvsmount	100044	MVSmount daemon (for **mvslogin**, **mvslogout**)
showattrd	100059	**showattr** daemon
pcnfsd	150001	**pcnfs** daemon
nlm	100021	Network Lock Manager
nsm	100024	Network Status Monitor

Setting up a user specified port range

The /etc/rpc file for the z/OS NFS client must define entries for services **mvsmount** and **showattr**. Table 28 outlines the services number required for **mvsmount** and **showattr**.

Table 28. Modify /etc/rpc for mvsmount and showattr

Service	Number	Description
mvsmount	100044	MVSmount daemon (for **mvslogin**, **mvslogout**)
showattrd	100059	**showattr** daemon

The /etc/services file for the z/OS NFS server must define the port number entries for services **mountd**, **mvsmount**, **pcnfsd**, and **showattr**. Figure 20 outlines the port numbers for these services, with contiguous port numbers 2045-2048 as examples.

```
#
# NFS server
#
#  Port 2049 is an exclusive port for nfsd and should not be
#       changed.
#  Ports 2045-2048 are interchangable for mountd, mvsmount,
#       showattr, and pcnfsd.
#
```

Figure 20. Modify /etc/services for mountd, mvsmount, pcnfsd, and showattr (Part 1 of 2)

```
status      2043/tcp   nfs_state # NFS State daemon (NLM)
status      2043/udp              # NFS State daemon (NLM)
status      2043/udp              # NFS State daemon (NLM)
nlockmgr    2044/tcp   nfs_lockd # NFS Lock daemon (NLM)
nlockmgr    2044/udp              # NFS Lock daemon (NLM)
mountd      2045/tcp   nfs_mount # NFS mount daemon
mountd      2045/udp              # NFS mount daemon
mvsmount    2046/tcp   nfs_mvsm  # NFS mvsmount daemon
mvsmount    2046/udp              # NFS mvsmount daemon
showattr    2047/tcp   nfs_showa # NFS showattr daemon
showattr    2047/udp              # NFS showattr daemon
pcnfsd      2048/udp   nfs_pcnfs # NFS pcnfsd daemon
nfsd        2049/tcp   nfs       # NFS server daemon
                                  #  - do not change
nfsd        2049/udp   nfs       # NFS server daemon
                                  #  - do not change
```

Figure 20. Modify /etc/services for mountd, mvsmount, pcnfsd, and showattr (Part 2 of 2)

The tcpip.profile file for the z/OS NFS server must define the port range entries for services **mountd**, **mvsmount**, **pcnfsd**, and **showattr**. Figure 21 outlines the port ranges for these four services with contiguous port numbers 2045-2048 as examples, starting with ports 2043 and 2044 for the network lock manager and network status monitor.

```
PORTRANGE 2043 7 UDP mvsnfs ; Reserved for startup JCL,
                            ;     mvsnfs
PORTRANGE 2043 7 TCP mvsnfs ; Reserved for startup JCL,
                            ;     mvsnfs
```

Figure 21. Modify tcpip.profile for NLM integration

Configuring a secure z/OS NFS server

In order for the z/OS NFS version 4 server to be able to provide RPCSEC_GSS security authentication flavors such as krb5, krb5i and krb5p, the z/OS NFS server needs to be configured to communicate with the Kerberos facilities. To do so, complete the following steps.

These steps assume that Resource Access Control Facility (RACF) is available in the system. If you have a different but equivalent external security manager, please refer to the product documentation for instructions.

1. The Kerberos key distribution center (KDC) must be running, and must contain the z/OS NFS server's principal before the z/OS NFS server starts. Otherwise, the following message will appear when z/OS NFS server starts:

   ```
   GFSA737I NETWORK FILE SYSTEM SERVER KERBEROS TICKET ACQUISITION FAILED
   ```

   In this case z/OS NFS server will be functional and only allow **authsys**, but not the other Kerberos authentication levels such as **krb5**, **krb5i**, and **krb5p**.

   The KDC can be running on z/OS, either on the same host as the z/OS NFS server itself or remotely from the z/OS NFS server. It can also be a KDC running on other platforms, for example, a SUN Solaris system or any other platform.

   For a brief description on how to setup a z/OS KDC, please see Appendix N, "Setting up a Kerberos Key Distribution Center," on page 391, or refer to *z/OS Integrated Security Services Network Authentication Service Administration*, SC24-5926 for more advanced details. For setting up other platforms KDCs, please refer to the specific platform's documentation.
2. Define local realm and default policy. For example, issue the following Ishell command:

   ```
   RDEFINE REALM KERBDFLT KERB(KERBNAME(REALM_NAME) PASSWORD(password)
   ```
3. Define IRR.RUSERMAP and grant READ authority to all system users, issuing Ishell commands:

   ```
   RDEFINE FACILITY IRR.RUSERMAP UACC(READ)
   SETROPTS RACLIST (FACILITY) REFRESH
   PERMIT IRR.RUSERMAP CLASS(FACILITY) ID(MVSNFS) ACCESS(READ)
   SETROPTS CLASSACT (FACILITY)
   ```
4. Create RACF user ids with Kerberos segments on the z/OS NFS server. For example, if a user is using a Kerberos principal "user1" on the NFS client, the Kerberos segment "user1" must be defined to a z/OS RACF user on the z/OS NFS server.

   For example: To add a RACF id, issue the Ishell command:

   ```
   AU (RACFID1) OWNER (IBMUSER) OMVS(UID(101))
   ```

   To add a Kerberos segment to the above user definition, issue the Ishell command:

   ```
   ALTUSER RACFID1 PASSWORD(password) NOEXPIRED KERB  (KERBNAME(user1))
   ```

   If the Kerberos segment is not defined correctly to RACF, the following error message appears on the server when a NFS client tries to mount to z/OS NFS server with Kerberos.

   ```
   GFSA728E SAF APPLICATION USER MAPPING FAILED WITH SAF RETURN CODE 8,
    RACF RETURN CODE 8, AND RACF REASON CODE 16.
   ```

Note: The ALTUSER command converts the password to upper case if the MIXEDCASE SETROPTS option is not set. If MIXEDCASE is not set, you must ensure that the uppercase value is used when you request an initial ticket. The principal name is not converted to upper case and the realm name is not included. You must change the password for the user in order to create the Kerberos secret key.

5. Add this path in the z/OS UNIX /.profile: "PATH=/usr/lpp/skrb/bin:$PATH" and export the "PATH".
6. Create or edit the krb5.conf file to contain the correct realm information for the z/OS NFS server. By default, the location of the krb5.conf file is in directory /etc/skrb.

   Sample /etc/skrb/krb5.conf file to be put on the z/OS NFS server host:

```
[libdefaults]
default_realm = KRB390.IBM.COM
kdc_default_options = 0x40000010
use_dns_lookup = 0
default_tkt_enctypes = des-cbc-crc
default_tgs_enctypes = des-cbc-crc

[realms]
KRB390.IBM.COM = {
 kdc = dcesec4.krb390.ibm.com:88
 kpasswd_server = dcesec4.krb390.ibm.com:464
 admin_server = dcesec4.krb390.ibm.com:749
}

KRB2000.IBM.COM = {
 kdc = sstone1.krb2000.ibm.com:88
 admin_server = sstone1.krb2000.ibm.com:749
}

[domain_realm]
.krb390.ibm.com = KRB390.IBM.COM
.krb2000.ibm.com = KRB2000.IBM.COM

[capaths]
KRB390.IBM.COM = {
  KRB2000.IBM.COM = .
}
```

   Note: As of July 2005, Linux (Enterprise Linux 4) and AIX (version 5.3) do not support the des-cbc-md5 encryption type. If the z/OS NFS server will be supporting multiple platforms of NFS client, IBM recommends using des-cbc-crc encryption types only, as shown above.
7. Generate the keytab from the KDC and put it in /etc/skrb of the z/OS NFS server unless otherwise defined. Examples of generating keytabs can be found in Appendix N, "Setting up a Kerberos Key Distribution Center," on page 391 or in *z/OS Integrated Security Services Network Authentication Service Administration*, SC24-5926.
8. If there is any multi-realm setup in the environment, the z/OS NFS server needs to have the foreign principals mapped to RACF. For example:

   To map a foreign principal "princ" in realm2 to RACF, issue Ishell command:

```
RDEFINE KERBLINK /.../REMOTE_REALM/princ  APPLDATA('RACF_ID')
```

   To map the entire foreign realm (every principal in the trusted foreign realm) to a RACF user, issue Ishell command:

```
RDEFINE KERBLINK /.../REMOTE_REALM/   APPLDATA('RACF_ID')
```

Note: the /.../ and trailing slash are required.

9. Start the z/OS NFS server. If set up is correct, the following message should be shown:

```
GFSA730I NETWORK FILE SYSTEM SERVER KERBEROS INITIALIZATION SUCCESSFUL
```

Notes:

1. The above are the minimal requirements to set up a secure z/OS NFS server in order for it to communicate with Kerberos facilities. For more advanced configurations, please see *z/OS Integrated Security Services Network Authentication Service Administration*, SC24-5926.
2. If z/OS NFS server is configured to use a KDC which resides on a remote host, the local KDC procedure (for example, skrbkdc) on the same host as the z/OS NFS server should not be started.
3. If the security site attribute is set to SAF or SAFEXP, in order to communicate with the z/OS NFS server using Kerberos, the MVS ID used by the mvslogin tool should be the RACF ID of the Kerberos segment that is used in the "kinit" client command. For instance, to access a SAF z/OS NFS server using Kerberos, and RACF has an user ID RACFID1, which has a Kerberos segment USER1, a user will need to perform "kinit USER1" on a NFS client to acquire the principal, and also issue command "mvslogin RACFID1" to get authorized by RACF, in order to be able to mount and access the SAF or SAFEXP enabled z/OS NFS server. In another example for the Kerberos multi-realm environment, if an entire trusted remote realm is mapped to RACFID2 with Kerberos segment USER2, then all users in the remote realm will have to perform "mvslogin RACFID2" and "kinit USER2" in order to access SAF or SAFEXP enabled z/OS NFS server.
4. For Linux (Enterprise Linux 4) users, as of July 2005, Linux NFS version 4 is still considered experimental. Although base NFS version 4 function is considered complete, RPCSEC_GSS is still under development. Linux Kerberos works differently from other UNIX platforms.
   - In order to perform a secure Kerberos mount, Linux requires a keytab generated by the Kerberos KDC to be put in the Linux machine's /etc. This behavior will change once kernel keyring support is completed.
   - Because Linux has the keytab, the system is able to perform secure NFS mount without having the credentials acquired by "kinit" command. This behavior will change once kernel keyring support is completed.
   - Root user (uid = 0) uses Linux machine credentials but not the regular user credentials obtained via kinit. Thus root user will be able to browse the NFS mount point without performing kinit. Regular users will need "kinit" to access the mount points. This behavior will change once kernel keyring support is completed.
   - Kdestroy will not destroy the context in the Linux kernel. This behavior will change once kernel keyring support is completed.
   - For Linux, multiple mounts to the same server with different security flavors are not supported yet.
   - Extra configurations are needed for Linux remote realm setup since Linux sends nfs/host.domain instead of user principal during mount time. If the Linux's NFS principal is not defined to RACF, z/OS NFS server will reject mount requests. A simple way to solve this problem is to map the entire remote realm to RACF. Another more secure way to work around this is to map individual Linux machines to a special realm in the [domain_realm] section in the /etc/skrb/krb5.conf, and map that realm to a special RACF on the z/OS NFS server thus leaving all other machines in the remote realm intact.

- For more information, please refer to http://www.citi.umich.edu/projects/nfsv4/linux/faq

5. For Windows users, it is recommended to use the MIT Leash Kerberos Ticket Manager with Hummingbird Maestro NFS client, especially when using Windows XP. However, to use Hummingbird NFS Maestro version 9 or below with Windows XP, a user might need to use Microsoft SSPI instead of the MIT Leash Kerberos Ticket Manager. For more information on Hummingbird NFS Maestro client and the configurations of the Kerberos on Windows platform, please consult Hummingbird NFS Maestro User's Guide.

   If the z/OS NFS server security site attribute is set to SAF or SAFEXP, the only supported authentication level on Windows platform is authsys. RPCSEC_GSS authentication levels such as krb5, krb5i and krb5p are not supported because Hummingbird Maestro 9 and 10 require NIS/NIS+ or LDAP with Kerberos instead of PCNFSD, which serves as the mvslogin tool for Windows platform. Thus, if z/OS NFS server is SAF enabled and the network transmission protection is mandatory, the Windows platform is not recommended.

Using dynamic client IP addressing

By default, the z/OS NFS server expects to communicate with clients based on a static client IP address. The server can also use the dynamic host configuration protocol (DHCP) to accept dynamic client IP address changes. To use dynamic client IP addressing, specify the **dhcp** server site attribute. The default attribute, **nodhcp**, tells the server to use the static IP algorithm.

To use dynamic IP addressing, the client must:

- Have a constant host name that the NFS server can identify it by.
- Dynamically update the authentication DNS (dynamic name server) with new IP addresses whenever they change.
- Maintain the TTL (time to live) value that the authentication DNS server specifies to any caching DNS server, based on the frequency with which system IP addresses might change. The larger the TTL value, the greater the possibility that the caching DNS server will have obsolete information. If dynamic addressing is used, the TTL value should be small, ideally zero, but a small value defeats the benefit of caching, so a compromise must be set with the understanding that cached values can become obsolete during the TTL interval and report incorrect information to querying systems like the NFS server.

Regardless of the **dhcp/nodhcp** attribute value, the z/OS NFS server itself continues to have a static IP address.

Installing the client enabling commands

This section describes the tasks you must perform to install and port the client enabling commands. These tasks include retrieving commands for AIX and Sun Solaris. This section also includes information about porting the **mvslogin**, **mvslogout**, and **showattr** commands and dealing with different compilers and operating systems.

To enable client users to access the z/OS system and to display system attributes, you must install the **mvslogin**, **mvslogout**, and **showattr** commands on the client workstations. For some client machines, you might need to modify the code to port these commands so they run on your client machine. See "Porting the mvslogin,

mvslogout, and showattr commands" on page 153. Before you install the commands, make sure that TCP/IP and File Transfer Protocol (FTP) are running both on z/OS and on the client.

Note: The z/OS NFS client utilities, including **mvslogin**, **mvslogout**, and **showattr**, are installed when the z/OS NFS client and TCP/IP are installed. The target library NFSCUTIL is a DDDEF to an existing z/OS UNIX directory (/usr/lpp/NFS/IBM) and will contain the client commands for the z/OS NFS client after installation. There is no need to port the z/OS NFS client utilities as you would for the remote NFS clients which use the z/OS NFS server.

Follow these installation procedures:

1. Delete any previous versions of the **mvslogin**, **showattr**, and **mvslogout** commands and their source code from your client workstation.
2. Retrieve the **mvslogin**, **mvslogout**, and **showattr** commands (in tarbin file format for AIX or UNIX, or in source code format for any other platform) from the *prefix*.NFSTARB data set, where prefix is an installation-specified variable. Use FTP with a binary transfer to send the tarbin or executable files to a client workstation for UNIX or AIX (no character conversion should be made). Use FTP with text transfer to send the source code of the three commands to a client workstation for any platform.
3. Use the tar utility to extract the files only if the client uses AIX or UNIX.
4. Compile the source code only if it is not in executable code format. (You might need to modify the code for your specific client machine.)
5. Make executable code versions of the commands available to all clients.

Recommendation: We recommend placing these commands on a LAN server (possibly in /usr/local/bin) that is available to many workstations, rather than installing them on each client workstation.

Table 29 is a list of all files stored in *prefix*.NFSTARB data set related to the server's client commands (**mvslogin**, **mvslogout**, and **showattr**):

Table 29. Files in the prefix.NFSTARB data set to download to clients

File Name	Download as:	Description	Client Environment
prefix.NFSTARB(GFSAWAIX)	client.tarbin (or any name)	Binary file	AIX, UNIX
Source Code for Commands			
prefix.NFSTARB(GFSAWMNT)	gfsawmnt.h	C header files	All [1]
prefix.NFSTARB(GFSAWSHO)	gfsawsho.h		
prefix.NFSTARB(GFSAWRP6)	gfsawrp6.h		
prefix.NFSTARB(GFSAWRS6)	gfsawrs6.h		
prefix.NFSTARB(GFSAWAXD)	gfsawaxd.c	C modules	All [1]
prefix.NFSTARB(GFSAWCLT)	gfsawclt.c		
prefix.NFSTARB(GFSAWLIN)	gfsawlin.c		
prefix.NFSTARB(GFSAWLOU)	gfsawlou.c		
prefix.NFSTARB(GFSAWMCL)	gfsawmcl.c		
prefix.NFSTARB(GFSAWMOU)	gfsawmou.c		
prefix.NFSTARB(GFSAWSHA)	gfsawsha.c		
prefix.NFSTARB(GFSAWMAK)	makefile	All [1]	

Note:
1. For AIX or UNIX, you do not need to download every individual file if you download the GFSAWAIX file.

The following sample client screens show how to retrieve and create the **mvslogin**, **mvslogout**, and **showattr** client commands for the following platforms.

- AIX, Sun Solaris, and Linux; see page 150.

Appendix I, "Retrieval of source code for client enabling commands," on page 359 shows how to retrieve the necessary source code to install the client commands on any platform except for an AIX or UNIX workstation.

Retrieving commands for AIX, Sun Solaris, and Linux

Figure 22 shows how to retrieve the necessary files to install the client commands on workstations with AIX, Sun Solaris, or Linux. Use the ″Make″ section that is appropriate for your platform and omit the others. Use MVSHOST1 as the name of the z/OS host, and smith is a z/OS user ID. Note: the compiler name can be changed to match the installed compiler name.

```
$ ftp mvshost1
Connected to mvshost1
220-FTPSERVE at MVSHOST1, 01:44:24 on 6/02/94
220 Connection closes if idle for more than 5 minutes.
Name (mvshost1:w42dept):  smith
<Press ENTER key>
331 Send password please.
Password: password
230 smith is logged on.
ftp> bin
200 Representation type is IMAGE.
ftp> get 'prefix.nfstarb(gfsawaix)' client.tarbin
200 Port request OK.
125 Sending data set PREFIX.NFSTARB(GFSAWAIX)
250 Transfer completed successfully.
local: client.tarbin remote: 'prefix.nfstarb(gfsawaix)'
213504 bytes received in 2.4 seconds (87 Kbytes/s)
ftp> quit
221 Quit command received. Goodbye.
$
$ mkdir mvsnfs.client
$ cd mvsnfs.client $ tar -xvf ../client.tarbin
tar: record size = 20 blocks
x ./makefile, 8234 bytes, 17 tape blocks
x ./gfsawaxd.c, 13474 bytes, 27 tape blocks

  ⋮
$
$ touch *.*
$ ftp mvshost1
Connected to mvshost1
220-FTPSERVE at MVSHOST1, 01:44:24 on 6/02/94
220 Connection closes if idle for more than 5 minutes.
Name (mvshost1:w42dept): smith
<Press ENTER key>
331 Send password please.
Password: password
230 smith is logged on.
ftp> bin
200 Representation type is IMAGE.
ftp> get 'prefix.nfstarb(gfsawaix)' client.tarbin
```

Figure 22. Retrieving the client enabling commands for AIX, Sun Solaris, and Linux (Part 1 of 2)

```
 200 Port request OK.
 125 Sending data set PREFIX.NFSTARB(GFSAWAIX)
 250 Transfer completed successfully.
 local: client.tarbin remote: 'prefix.nfstarb(gfsawaix)'
 213504 bytes received in 2.4 seconds (87 Kbytes/s)
 ftp> quit
 221 Quit command received. Goodbye.  $
 $ mkdir mvsnfs.client
 $ cd mvsnfs.client
 $ tar -xvf ../client.tarbin
 tar: record size = 20 blocks
 x ./makefile, 8234 bytes, 17 tape blocks
 x ./gfsawaxd.c, 13474 bytes, 27 tape blocks
   :
   :
 $
 $ touch *.*

(Following are Make sections for each supported platform - use the appropriate one

SUN:
$ make sun
gcc -D SUN -D _UTILS -D SOLARIS -lnsl -lsocket -o showattr gfsawsha.c gfsawaxd.c
gcc -D SUN -D _UTILS -D SOLARIS -lnsl -lsocket -o mvslogin gfsawlin.c gfsawmou.c
gcc -D SUN -D _UTILS -D SOLARIS -lnsl -lsocket -o mvslogout gfsawmou.c gfsawlou.c

Sun binaries for NFS Tools are made

LINUX:
make linux
gcc -D LINUX -D _UTILS -o showattr gfsawsha.c gfsawaxd.c
gcc -D LINUX -D _UTILS -o mvslogin gfsawlin.c gfsawmou.c
gcc -D LINUX -D _UTILS -o mvslogout gfsawmou.c gfsawlou.c

Linux binaries for NFS Tools are made

AIX:
$ make aix
cc -D AIX -D _UTILS -D AIX_RT -o showattr gfsawsha.c gfsawaxd.c
cc -D AIX -D _UTILS -D AIX_RT -o mvslogin gfsawlin.c gfsawmou.c
cc -D AIX -D _UTILS -D AIX_RT -o mvslogout gfsawmou.c gfsawlou.c

AIX binaries for NFS Tools are made
 $ ./mvslogin mvshost1 smith    (MVSNFS must be operational
 Password Required                          on host side.)
 GFSA973A  Enter MVS password: password
 GFSA955I  smith logged in ok.
 $
```

Figure 22. Retrieving the client enabling commands for AIX, Sun Solaris, and Linux (Part 2 of 2)

After retrieving the files of client enabling source code for the AIX or UNIX environment, follow these steps as shown in Figure 22 on page 150 to create the executable commands:

1. To make sure the source files have the current date. Some platforms do not have correct time stamps and cause make files to fail; issue the following command.

   ```
   touch *.*
   ```

2. To create the executable commands **mvslogin**, **mvslogout**, and **showattr**, issue one of these commands for the environment that IBM supports.

   make aix
   Creates executables for AIX.

make sun
Creates executables for Sun.

make linux
Creates executables for Linux.

Porting the mvslogin, mvslogout, and showattr commands

The z/OS NFS server supports any client machine that has an NFS client software package implemented according to the Sun NFS protocol. The z/OS NFS server has been tested with the following client platforms:

- AIX 5.3
- Sun Solaris 10
- Windows Hummingbird Maestro 9
- Redhat Enterprise Linux 4
- z/OS NFS

Note: NFS supports the authentication procedures of PCNFSD Version 1 and Version 2 protocols. If a PC client supports PCNFSD and keeps the UID and GID to each mount point base, you do not need to port the **mvslogin** command. See Appendix J, "PCNFSD protocol," on page 361 for details on PCNFSD support.

To port the **mvslogin**, **mvslogout**, and **showattr** commands successfully, you should understand the following:

- C language - The source code for these commands is written in C.
- System calls for your client machine's operating system - For example, the FAT file system under DOS only allows up to eight characters for file names, and up to three characters for file name extensions. AIX and UNIX do not have this restriction. Therefore, while mvslogout is a valid file name in an AIX or UNIX environment, it is too long to be a valid file name in a FAT file system under DOS.

  As another example, the way that you get mount information varies for different platforms. The **mount** command is in the following (or similar) format:

  ```
  mount <server>:<remote file system> <local mount point>
  ```

  The minimum information for porting the client enabling commands is:

  1. Server name
  2. Remote file system (high-level qualifier)
  3. Local Mount point
  4. UID and GID

  The system calls to get the information for porting the client enabling commands are platform-dependent. If you cannot find the information in the following types of documents for the platform, you must call the support telephone number for the platform and ask to speak with their NFS development department:

  – Operating system development toolkit
  – TCP/IP development toolkit
  – NFS development toolkit
- The source code for **mvslogin**, **mvslogout**, and **showattr**.

  For example, mvslogin tells the server the z/OS user id and its associated client UID number. This client UID number is expected to be passed to the server for all further client requests to the NFS. If the client user does not specify the z/OS user id and password on the mvslogin command, the z/OS login ID is taken from the login ID on the workstation with no password assumed. If authentication for this default login ID from the workstation fails, then mvslogin prompts the user to enter the z/OS login password.

Figure 23 shows the common source files for the **mvslogin**, **mvslogout**, and **showattr** commands on all platforms:

Seven .c files

gfsawaxd.c
XDR encode and decode routines for attributes service.

gfsawclt.c
NFS protocol calls for **mvslogin** and **mvslogout**.

gfsawlin.c
Main program to generate **mvslogin** command.

gfsawlou.c
Main program to generate **mvslogout** command.

gfsawmcl.c
Create the client handle and initialize it.

gfsawmou.c
XDR protocol definitions for **mvslogin** and **mvslogout**.

gfsawsha.c
Main program to generate **showattr** command **mvslogout**.

Four .h files

gfsawmnt.h
Protocol definitions for **mvslogin** and **mvslogout**.

gfsawsho.h
Attribute definition and procedures.

gfsawrp6.h
TCP/IP version 6 RPC library definitions.

gfsawrs6.h
TCP/IP version 6 RPC support functions.

One makefile

Figure 23. Common source files

Porting on different compilers and operating systems

Procedures for porting vary for different C compilers and operating systems. Differences can occur during compiling, linking, and run time.

Compiling

The following items might vary for your client machine's operating system:

- Different set of compilation flags

  There are different sets of compilation flags based on compilers or operating systems. For example:
  - AIX (on RS/6000) has the unique flags _BSD, _SUN, and BSD_INCLUDES.
  - DOS compilers have different compiler models, which require the corresponding compiler flag (for example, -AL and -AS).
- Include files in different directories

  Because the include files can be installed differently based on the operating systems and their toolkits, the include files could be in different directories.
- Include file has a different name

Include files for the same or similar functions could have the same or similar file names. For example, DOS uses the file name “string.h”, and the other platforms use “strings.h”.

- System variables in different include files

  The system variables are usually in different include files, based on the operating systems. For example, to access the mount table some AIX and UNIX clients use mntent.h.

- System variables have different names

  The variables related to operating systems could have different variable names. For example; some AIX and UNIX clients use getuid to get the real UID.

- System variables have different structure

  The structures related to operating systems could be different. For example, DOS FAT file systems have file name length restrictions which cause them to have a different directory structure from AIX or UNIX.

- System variables not supported

  Some system variables are supported by one operating system but not another.

- Sequence of include files

  Some include files are based on the precedence of another include file. The prerequisite include file has to come before the other include files. For example, some Programming Libraries offer types.h which is based on C compiler sys/types.h. Therefore, #include <sys/types.h> should be before #include <types.h>.

- Mount information varies depending on the client operating system

  The information about mount points provided by vendors of the client operating systems and client TCP/IP products varies and might not always be complete. To find the mount information:

  1. Search through the documentation (for example, the TCP/IP development toolkit and the installation and administration guides).
  2. If you cannot find the mount information in the documentation, contact the vendor that offers the TCP/IP development toolkit.

Linking

After linking the programs together, check for attention messages and error messages. The following items might vary for your client machine’s operating system:

- Different set of linker/loader flags

  Some programs require a different set of linker/loader flags.

- Library files in a different directory

  The library files required to complete the linkage could be in a different directory.

- Library files have different names

  Depending on how the client machines’ operating systems are installed, the library files might have different names.

- Different libraries required

  The system variables could be in different libraries for the different operating systems.

- Compiler is not compatible with the system toolkit

  Some operating systems support multiple versions of C compilers. Some C compilers, however, might not match the various system toolkits.

- Different library model required

The library models have to match with the compilation time.

- System variables not supported

  Some system variables are supported during compile time but not supported by the link time.

Running commands

After the compilation and linkage are successful, run the command to see if the result is as expected. If not, figure out the difference of the result or failure. The difference or failure can be in the following areas:

- Definition of the standard C variables is different

  The definition of the standard C variables could be different for the different operating systems or compilers. Some special handling might be required. For example:

  – **int** is 2 bytes long for DOS but it is 4 bytes long for the z/OS NFS, AIX, and Sun.

  – **tab** has a different value, causing the spacing of the output to be different.

- Definition of function calls is slightly different

  Although a given function is supported, it might work slightly differently. For example, the "mount table" has a different format in AIX from Sun Solaris.

- Library functions might have a defect

  Some functions in the C libraries do not function as the documentation describes. You might report the problems or write your own functions to replace them.

Chapter 11. Network File System operation

This chapter describes how to start and stop the z/OS Network File System, and describes the operands of the **modify** command that are related to the z/OS NFS server. The operands for collecting diagnostic information are also described.

Starting the z/OS NFS client

If you want to use the z/OS NFS client, FMID HDZ11VC, do the following:

- Define the z/OS client as a file system in the z/OS UNIX BPXPRMxx member of SYS1.PARMLIB. Start the z/OS UNIX address space, which will cause the BPXPRMxx member to be used. As part of the z/OS UNIX startup, the z/OS NFS client will be started in an z/OS UNIX colony address space. Wait until this message appears before proceeding:

  ```
  BPXI004I OMVS INITIALIZATION COMPLETE
  ```

  See *z/OS UNIX System Services Planning* for more information.
- Ensure that TCP/IP and PORTMAP are active.

During z/OS UNIX file system initialization, the z/OS NFS client is started and run in the LFS colony address space. The FILESYSTYPE parmlib statement for the z/OS NFS client must be present in the SYS1.PARMLIB(BPXPRMxx) parmlib member in order to start the z/OS NFS client. For more information on z/OS UNIX file system, see *z/OS UNIX System Services File System Interface Reference*.

If the z/OS NFS client fails to initialize, a write to operator (WTO) message is issued to the operator console. The following conditions can cause z/OS NFS client initialization to fail.

- The z/OS NFS client is not started in a standalone colony address space.
- The z/OS NFS client is already started. Multiple instances of the z/OS NFS client on a single z/OS system is not supported.
- Using a security product that is downlevel, the z/OS NFS client requires RACF 2.2 or later.
- An incorrect parameter is specified in the installation parameters.
- Unicode services is not installed or is not active.

When the z/OS NFS client initializes, messages like these example messages are displayed on the operator's console.

```
$HASP373 NFSCR    STARTED
BPXI004I OMVS INITIALIZATION COMPLETE
GFSC700I z/OS NETWORK FILE SYSTEM CLIENT
(HDZ11VC) STARTED
```

If the z/OS NFS client is stopped, canceled, or for any other reason terminates, z/OS UNIX issues the following message:

```
BPXF032D FILESYSTYPE type TERMINATED. REPLY 'R' WHEN READY TO RESTART.
         REPLY 'I' TO IGNORE
```

To restart the z/OS NFS client, specify 'R' in reply to the message. Replying with 'I' will cause z/OS UNIX to ignore the termination of the z/OS NFS client. If 'I' was replied and you still wish to restart the z/OS NFS client, use the SET

© Copyright IBM Corp. 1991, 2006

command as follows. If z/OS UNIX has been initialized but the z/OS NFS client is not active, issue a SET OMVS= command to the BPXPRMxx member of SYS1.PARMLIB that defines the z/OS NFS file system. z/OS UNIX will then start the z/OS NFS client.

Stopping the z/OS NFS client

The z/OS NFS client is started when the z/OS UNIX file system is initialized and is persistent until z/OS UNIX is stopped. To stop the z/OS NFS client gracefully the system operator can use the modify operator command **stop** with the z/OS NFS client address space name; for example:

```
stop mvsnfsc
```

If the stop command fails to gracefully shut down the z/OS NFS client, the operator can force an abnormal termination using the operator command **cancel** with the z/OS NFS client address space name; for example:

```
cancel  mvsnfsc
```

It is imperative and necessary to stop the z/OS NFS client gracefully so it can save its important data (RPC transaction ID) for the subsequent restart.

Attention: The destruction of the z/OS NFS client address space can cause unpredictable abnormal z/OS UNIX address space behavior. If a z/OS UNIX process tries to access the NFS client data during its address space destruction, then an OC4 protection exception in the z/OS UNIX BPXVCLNY load module can occur.

To bring down the NFS client during shutdown, follow these steps:

Step 1. Cancel all colony address spaces

Step 2. The NFS client runs in a colony address space; to terminate it, you can enter either:

```
STOP     <nfsc>  .....where <nfsc> is the nfs client proc name
```

or

```
CANCEL   <nfsc>  .....if STOP did not work
```

Step 3. Terminate UNIX System Services

You can use the BPXSTOP tool; select UNIX tools from the z/OS UNIX System Services Tools & Toys web site.

```
http://www.ibm.com/servers/eserver/zseries/zos/unix/bpxa1toy.html
```

or

```
F BPXOINIT,SHUTDOWN=FORKINIT
```

Step 4. Stop JES2

Starting the z/OS NFS server

Make sure that z/OS UNIX is customized to be able to start automatically during IPL.

If you want to use the z/OS NFS server, FMID HDZ11VS, TCP/IP and PORTMAP need to be started and active on your system. Then, start the z/OS NFS server.

You might also need to start up NAMESRV to map machine names into their corresponding Internet addresses. If you use the dynamic host configuration protocol (DHCP), NAMESRV is required.

Note: PORTMAP is synonymous with portmapper which is a program provided by z/OS Communications Server that maps registered server programs to port numbers. NAMESRV is the cataloged procedure of the Domain Name Server which is provided by Communications Server that maps a host name to an internet address or an internet address to a host name. See *z/OS Communications Server: IP Configuration Guide* and *z/OS Communications Server: IP Configuration Reference*for information on configuring PORTMAP, starting PORTMAP, configuring the Domain Name Server, and starting NAMESRV automatically with z/OS Communications Server.

The z/OS NFS server does not support file persistence. That is, when the server is restarted, files cannot be accessed using old file handles.

To start the z/OS NFS server, enter the **start** command from a console. Enter the command as follows:

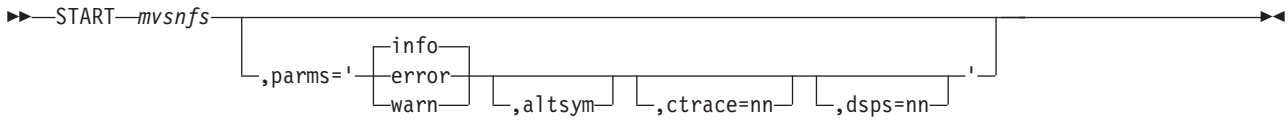

where

mvsnfs Specifies the name of the procedure in your system that PROCLIB used to start up the server.

info The first parameter specifies the level for diagnostic messages (the default is *info*).

altsym If *altsym* is specified in the second parameter, the semicolon (;) is used as the comment symbol in the attributes and exports data sets. Otherwise, the pound sign (#) is used as the comment symbol.

ctrace=nn
Specifies that diagnostic information for the NFS server be recorded in z/OS component trace buffers, using trace options specified in member CTINFS*nn* of SYS1.PARMLIB. To use the default trace options for the z/OS NFS server, specify *ctrace=00* for the default SYS1.PARMLIB member CTINFS00.

dsps=nn
Specifies the size of the data space to be allocated for the NFS component trace buffers, where (*nn* * 3) equals the number of megabytes to be allocated. The value range for *nn* is from 1 (the default), which allocates a 3 MB data space, to 600 (which allocates a 1.8 GB data space).

For example, suppose you have some data sets with a high-level qualifier of #*USER05*, and you want client users to be able to read them. First, you would modify the exports data set and attributes data set by using ';' as the comment symbol rather than '#'. Next you would specify `#user05 -ro` as an entry in the exports data set. Finally, you would specify the **altsym** parameter when you enter the **start** command.

These parameters override the parameter settings in the server startup procedure.

When you enter **start**, the following console message appears; if you installed HDZ11VS, then this FMID is displayed in the GFSA348I message text.

```
GFSA348I z/OS NETWORK FILE SYSTEM SERVER
(HDZ11VS) STARTED.
```

Starting multiple servers

You can start more than one z/OS NFS server on a single z/OS system. However, the z/OS V1R7 server (HDZ11VS) cannot coexist with servers started at a previous release level. If a V1R7 server is started first, then attempts to start previous release levels of the server will fail. If a previous-release server is running first, then a V1R7 level of the server cannot be started. Note: the NFS server and NFS client can be at different release levels on the same system – there is no requirement for server and client release levels to match.

Within a sysplex, you can have different levels of NFS servers and clients on different systems. However, byte-range locking and file share reservations are not communicated across systems within the sysplex. Therefore, in sysplex environments where byte range locking and share reservations are required for accessing MVS data sets via NFS, IBM recommends that only one NFS server be started.

Stopping the z/OS NFS server

When shutting down your system, follow these sequential steps.

Step 1. Stop the z/OS NFS server

Step 2. Shut down the TCP/IP server

Step 3. Shut down z/OS UNIX if it is running

Use the **stop** command to shut down the z/OS NFS server. All current input/output (I/O) operations are completed and all OPEN data sets are closed. This command can be entered at any time. Enter the command as follows, where *mvsnfs* is the name of the procedure in your system that PROCLIB used to start up the server.

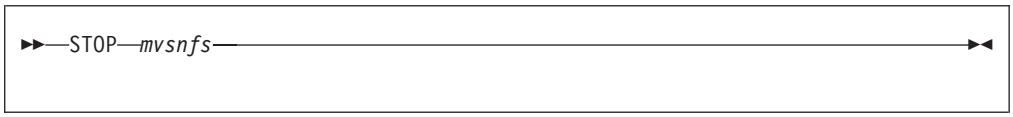

You can also shut down the server by entering the **modify** command.

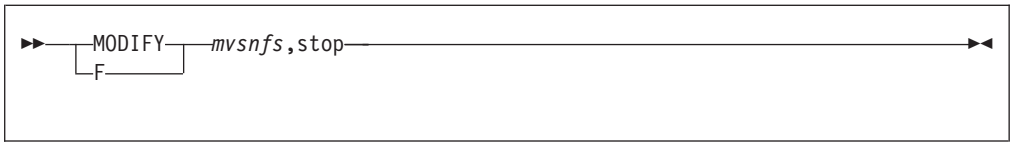

The operator's console displays messages like the following messages.

```
GFSA329I SERVER SHUTDOWN IN PROGRESS.
GFSA330I SERVER SHUTDOWN COMPLETE.
$HASP395 MVSNFS   ENDED
```

As a last resort, you can cancel the server by entering the **cancel** command.

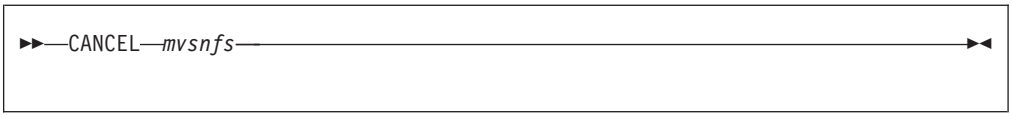

Starting the z/OS NFS NSM and z/OS NFS NLM

If you want to use the z/OS NFS Network Status Monitor (z/OS NFS NSM) and the z/OS NFS Network Lock Manager (z/OS NFS NLM), you must specify the NLM site attribute for the z/OS NFS server. See Table 23 on page 104 for details. If you specify the NLM site attribute, NLM and NSM will start when the NFS server is started, and will stop when the NFS server is stopped. The attribute's default value, NONLM, specifies that the NFS server will run without NLM or NSM.

Starting component tracing for the z/OS NFS server

To start recording diagnostic information for the z/OS NFS server in z/OS component trace buffers, follow these steps:

1. Decide on the trace options to use. These can be in a CTINFS*nn* member of SYS1.PARMLIB to be specified on the TRACE CT command, or individual options to be specified when prompted in response to the TRACE CT command.
2. From the master console or another console with master authority, issue the TRACE CT command as follows:

where

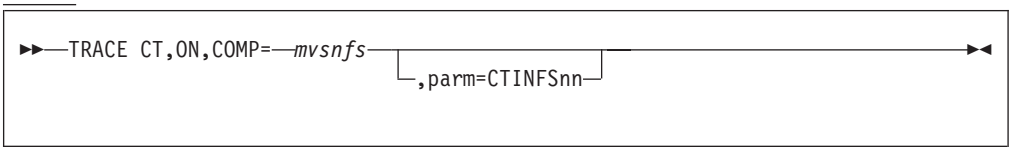

mvsnfs Specifies the name of the procedure in your system that PROCLIB used to start up the server.

parm=CTINFSnn
Specifies that diagnostic information for the NFS server be recorded in z/OS component trace buffers using trace options specified in member CTINFS*nn* of SYS1.PARMLIB. To use the default trace options for the z/OS NFS server, specify *parm=CTINFS00* for the default SYS1.PARMLIB member CTINFS00.

If you do not specify a CTINFS*nn* PARMLIB member on the CTRACE CT command, the following message will prompt you to enter trace operands:

```
* id ITT006A SPECIFY OPERAND(S) FOR TRACE CT COMMAND.
```

In response to the message, reply with the identification number *id* from the message, an external writer program if any to receive the records, and a trace option or options to use from the list below. Here is the response syntax:

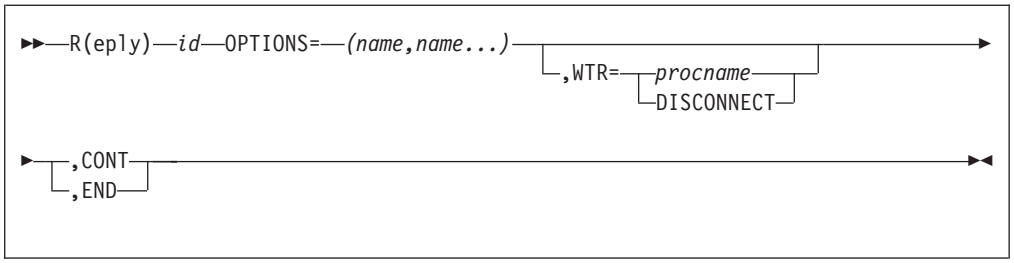

where:

id Specifies the identification number from the ITT006A message.

name,name...
Specifies a trace option or options for tracing NFS server records. The possible options are:

FFDC
First Failure Data Capture. This option is on by default.

Entry
entry into a function

Exit
exit from a function.

Suspend
unit of work must suspend due to unavailability of resource (e.g. waiting for a latch).

Resume
unit of work resumed due to availability of resource (e.g. latch acquired).

Schedule
unit of work is scheduled onto another queue (e.g. ipcqueue of this or another task, array queue, etc.).

Dispatch
unit of work is dispatched from a queue to resume processing.

USS_request
request issued to z/OS UNIX System Services.

USS_return
returned from z/OS UNIX System Services request.

DFP_request
Call to MVS data set system function.

DFP_return
Return from MVS data set system function.

CB_Mgmt
Control Block Management (i.e. creation, initialization, modification, deletion).

Network
Network communication related trace.

General
General Trace Record.

Detail
Detailed Trace Record. This is used for low level debug.

Trap
For use in special temporary trap code created to aid in the analysis of a problem.

Error
Existing NFS Trace Error Record.

Warning
Existing NFS Trace Warning Record.

Info
Existing NFS Trace Info Record.

Debug*x*
Existing NFS Trace Debug 1 - 9 Record.

In addition to the basic options above, you can also enter the following shorthand values to specify multiple record types:

All
All existing NFS trace records.

Call
Entry and Exit record types.

Task_Flow
Suspend, Resume, Schedule and Dispatch record types.

USS
USS_Request and USS_Return record types.

MVS
DFP_Request and DFP_Return record types.

WTR=procname | DISCONNECT
Connects or disconnects the component trace external writer and the trace. *procname* identifies the name of the member that contains the source JCL that invokes the external writer. The member can be a SYS1.PROCLIB cataloged procedure or a job. The procname in the WTR parameter must match the procname in a previous TRACE CT,WTRSTART command.

WTR=DISCONNECT disconnects the writer and the trace. The component continues tracing and placing the trace records in the address-space buffer, but stops passing trace records to the external writer.

You must also specify a TRACE CT,WTRSTART or TRACE CT,WTRSTOP command to start or stop the writer.

CONT or END
Specifies that the reply continues on another line. The system reissues the same prompting message. You then can continue the reply. You can repeat any parameters on the continuation line, except END. Repeated parameters are strung together. They do not overlay each other. You must specify END to complete the response. END identifies the end of the REPLY.

Note: Be sure to use the **flushctr** option on the Modify command to flush the last buffer being filled to the component trace before the external writer is stopped or before a dump is taken. Otherwise the records in that buffer won't be included in the component trace.

Entering operands of the modify command for the z/OS NFS server

Besides the **start** and **stop** commands, you can enter the **modify** command with these operands at the console: *exportfs*, *flushctr*, *freeze*, *list*, *release*, *status*, *unmount*, *flushlog*, *listlock*, *log*, *smf*, and *version*. The *flushctr*, *flushlog*, *listlock*, *log*, *smf*, and *version* operands can be used to collect data when diagnosing problems. All operands must be preceded by either `MODIFY` *mvsnfs*, or `F` *mvsnfs*, where *mvsnfs* is the name of the procedure in your system that PROCLIB used to start the server.

Each message is sent to a component trace buffer, or to the console or the data set that is pointed to by the NFSLOG1 and NFSLOG2 DD statements, or to any combination of those destinations. NFSLOG1 and NFSLOG2 are DD statements in the startup procedure of the z/OS NFS server. These z/OS NFS server DD statements specify data sets where all the messages for debugging or trace are stored during the processing of the z/OS NFS server.

All data set names entered from the console must be fully-qualified and without quotation marks.

This is a summary of the modify command with the relevant operands:

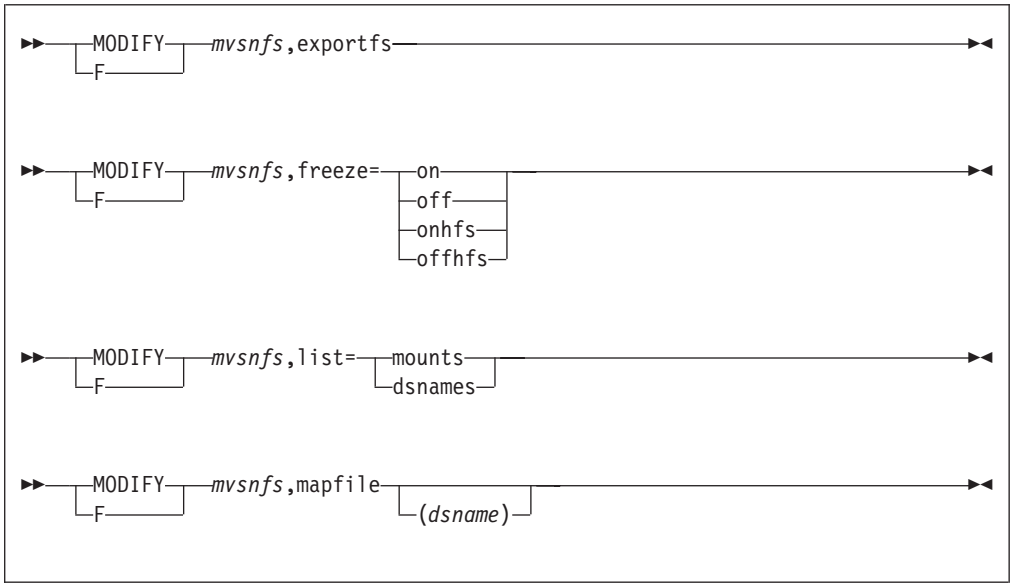

Figure 24. Summary of the modify command (Part 1 of 2)

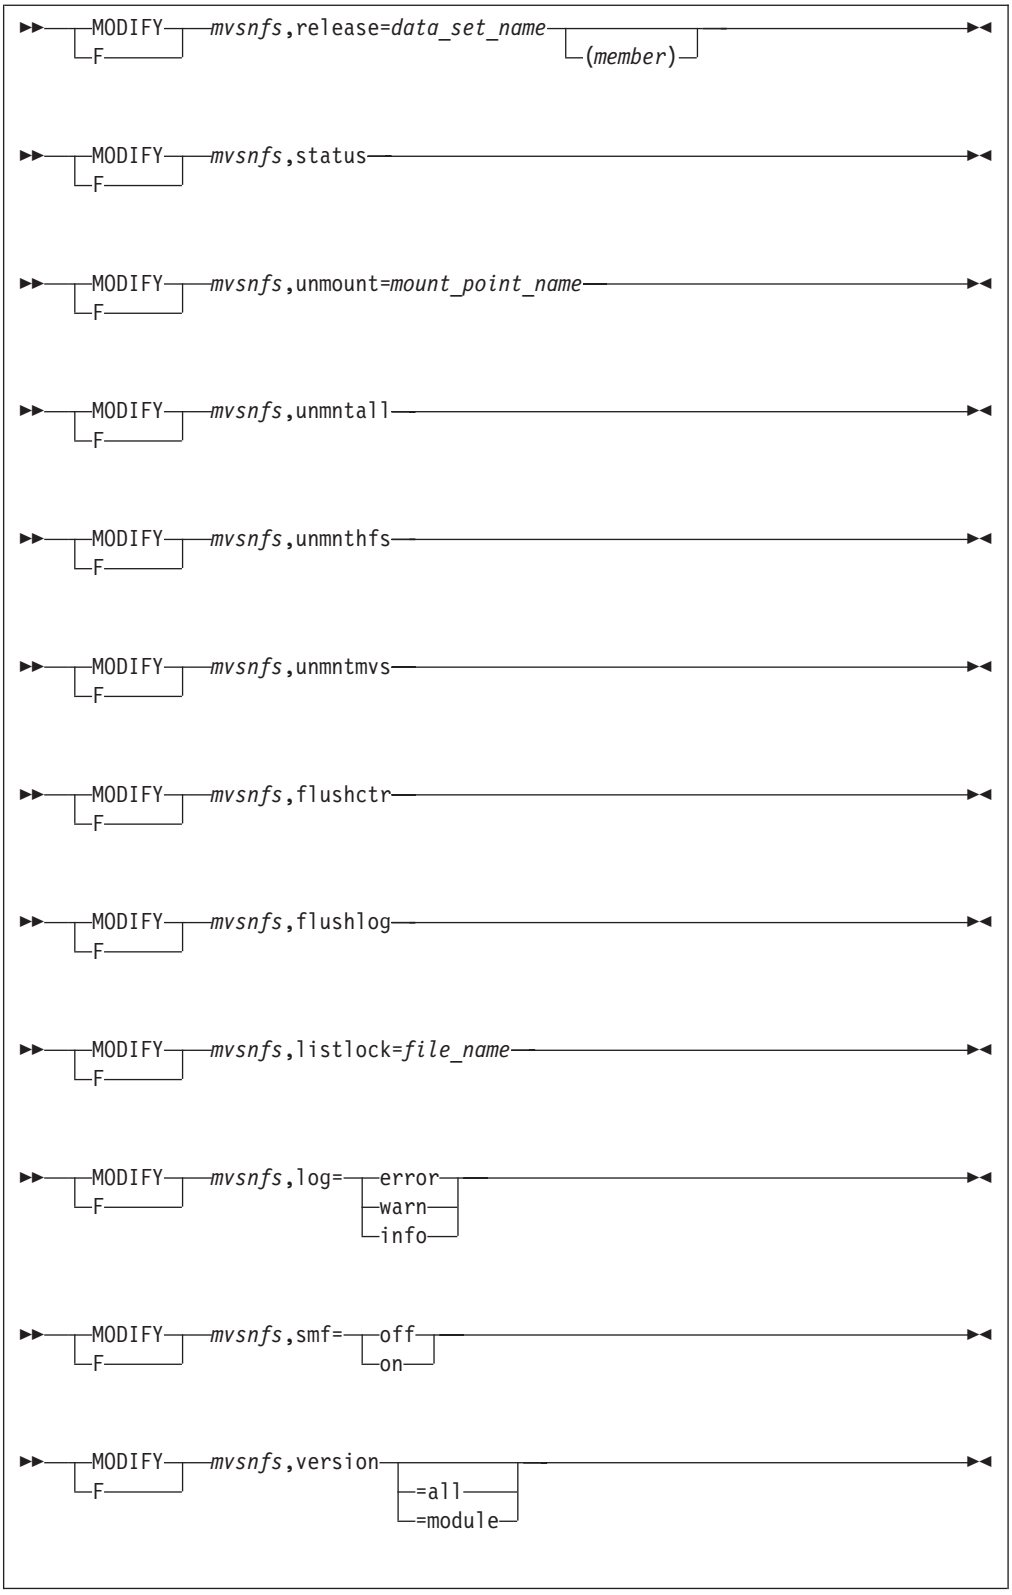

Figure 24. Summary of the modify command (Part 2 of 2)

Exportfs operand

This operand causes the exports data set to be re-read and the internal exports list to be rebuilt without stopping and restarting the server. This operand can be run at any time. Any changes in the **sec** keyword of the export list will be IMMEDIATELY enforced at server startup or EXPORTFS time, regardless of the state of any preexisting mount points. Enter the operand as follows.

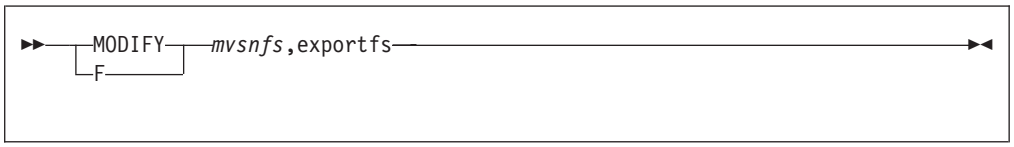

Freeze operand

This operand suspends or resumes processing of user mount requests. For the NFS V4 protocol, this operand also suspends or resumes processing of lookup operations from the NFS server into z/OS UNIX file systems or MVS data sets. You can enter the *freeze=on* command at any time for maintenance purposes. Enter the operand as follows.

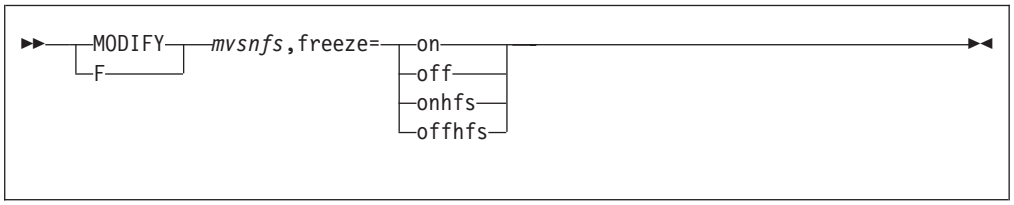

- If you select **on**, these messages appear.

  ```
  GFSA901I  MOUNT PROCESSING SUSPENDED.
  GFSA771I  mvsnfs HFS MOUNTS SUSPENDED.
  ```

  Future mount requests by client users are refused for both z/OS UNIX file systems and conventional MVS data sets and the message "Permission Denied" displays on their monitors. After issuing a *freeze=off* operand, mount processing resumes for both z/OS UNIX file systems and conventional MVS data sets, and these messages appear.

  ```
  GFSA902I  MOUNT PROCESSING RESUMED.
  GFSA772I  mvsnfs HFS MOUNTS RESUMED.
  ```

  Client users can again mount z/OS directories as normal.
- If you select **onhfs**, future mount requests by client users for z/OS UNIX file systems are refused.
- If you select **offhfs**, mount processing resumes for z/OS UNIX file systems.

List operand

This operand displays a list of either the mount points or the MVS data sets that are currently active in the Network File System. This command can be entered at any time for maintenance or diagnostic purposes. The path name output of an z/OS UNIX file object might require multiple console messages. Enter the operand as follows.

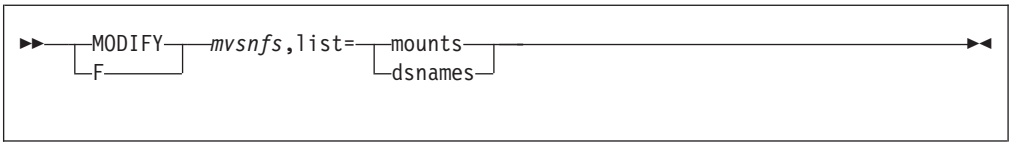

The **mounts** option returns a list of all mount points currently active in the system. The list of active mount point names and their associated *current use counts* are displayed as follows:

```
GFSA910I JOHN ACTIVE =12.
GFSA910I JOAN ACTIVE =10.
GFSA910I /HFS/U/BLUE ACTIVE =1.
```

Current use counts indicate how many mount requests have been made without an unmount request for the same mount point regardless to which local directory the mount is attached. For example, suppose the high-level qualifier `JOHN` is mounted to the same local directory 12 times without any unmount. `ACTIVE =12` is shown. Now, suppose the high-level qualifier `JOAN` has been mounted to 15 different local directories but 5 of them have been unmounted. `ACTIVE =10` is shown. With the NFS version 4 protocol, the current count reflects only mounts specified with **mvsmnt** , and the count is not decremented since NFS version 4 has no unmount function.

The **dsnames** option returns a list of all data sets that are currently active. Currently active means that the data set is either in use or has been opened for accessing but has not been closed due to timeout even though the file is not in use. A list of data sets and data set members such as these display.

```
GFSA912I CHRIS.TEST.
GFSA912I SMITH.PAYROLL(JULY).
GFSA912I /HFS/U/PAYROLL.
```

z/OS UNIX file names are truncated after 126 characters.

Mapfile operand

This operand causes the side files data to be read again and rebuilt without stopping and restarting the server or remounting of mount points. If mapfile operand has a parameter (particular side file name), only the specified side file will be updated.

The default is to refresh all side files (default from NFSATTR or previously specified by any mount command). The general command syntax follows.

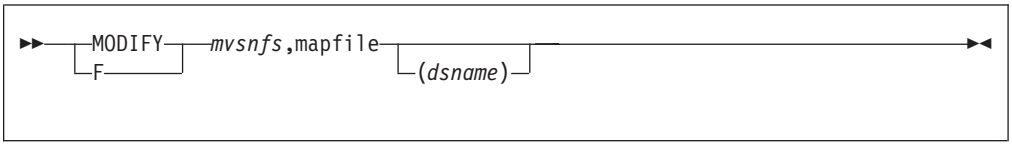

Release operand

This operand forces the Network File System to close and deallocate the MVS data set or PDS/PDSE member specified. Any locks which exist will also be released and any waiting locks for the data set or member will be cancelled. Any waiting byte range lock requests (for the specified data set) will be purged. File share reservations will also be released. For a z/OS UNIX file, any locks held will be released and any waiting locks will be canceled. z/OS UNIX file names must be

specified in single quotes with the hfs prefix (/hfs) followed by the z/OS UNIX path. This command can be entered at any time. The command syntax is as follows.

Warning: the NLM protocol does not provide any means for the NFS server to notify the NFS client that the locks have been released. Use of the RELEASE command against files with NLM locks can lead to data corruption because the NFS client will be unaware of the fact that its locks are no longer held.

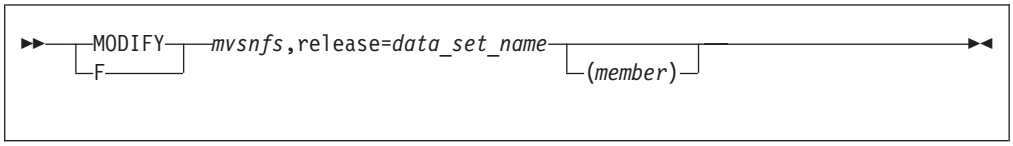

where

data_set_name
Is the name of an MVS data set or an z/OS UNIX file. If *data_set_name* references a PDS or a PDSE data set (does not include the member name), a syntax error will result.

For z/OS UNIX file names, specify the release operand as follows:

```
f mvsnfs,release='/hfs/u/jones' (hfs path specified in quotes)
```

For MVS data sets, specify the release operand as follows:

```
f mvsnfs,release=data_set_name (MVS dataset name not specified in quotes)
```

- If the object you specify is active, you receive a message as follows.

  ```
  GFSA914I data.set.name(member) DEALLOCATED.
  ```

- If the object you specify is not active, you receive a message as follows.

  ```
  GFSA915I data.set.name(member) NOT ALLOCATED.
  ```

Status operand

This operand displays the status of the server's active subtasks. You can enter this command for diagnosis purposes at any time that the server is running. Enter the operand as follows.

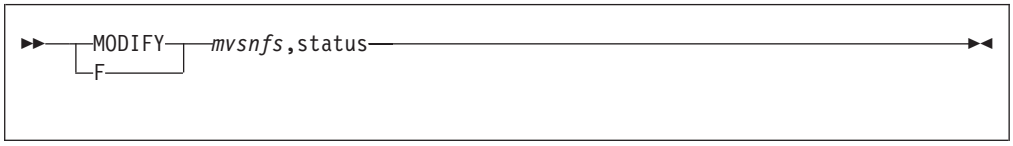

This is a sample message listing.

```
   GFSA900I MOUNT PROCESSING ACTIVE.
   GFSA751I SMF PROCESSING SUSPENDED FOR USER LOGOUT.
   GFSA753I SMF PROCESSING SUSPENDED FOR FILE TIMEOUT.
   GFSA781I mvsnfs SAF PROCESSING ENABLED.
   GFSA903I TASK 5C580 TCB 8D1888 PROGRAM = GFSALEGT =
NFSTSK02
   GFSA903I TASK 5C7A0 TCB 8D10C8 PROGRAM = GFSALEGT =
NFSTSK01
   GFSA903I TASK 5C9C0 TCB 8D1378 PROGRAM = GFSALEGT =
NFSTSK00
   GFSA903I TASK 5CE00 TCB 8D1378 PROGRAM = GFSAXPRT =
TRANSPORT
```

Unmount operand

This operand unmounts a mount point that is currently active. The path name for a z/OS UNIX file object cannot exceed 126 bytes. The `mount_point_name` for a z/OS UNIX path must be specified in single quotes with the hfs prefix (e.g. /HFS) in uppercase followed by the z/OS UNIX path (case sensitive). For example: `F mvsnfs,unmount='/HFS/u/jones'`. The general command syntax is as follows.

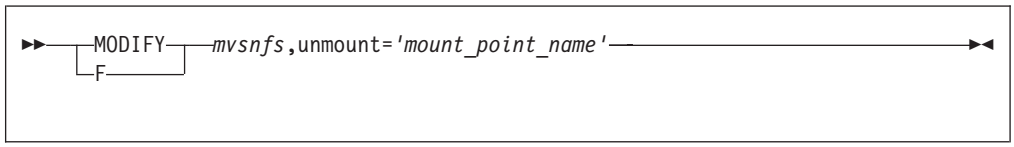

- If the mount point is active, the server responds.

  ```
  GFSA916I mount_point UNMOUNTED.
  ```

- If the mount point you specify does not exist when you enter the command, the server displays messages similar to the following messages.

  ```
  GFSA917I SMITH NOT MOUNTED.
  GFSA917I /HFS/U/JONES NOT MOUNTED.
  ```

Once the mount point is removed, client users are unable to access this mount point, and they get the "Stale NFS File Handle" message. Client users have to enter the **umount** command to end the stale file handle problem.

Unmntall operand

This operand causes the server to immediately unmount all mount points without stopping and restarting the NFS server. The general command syntax follows.

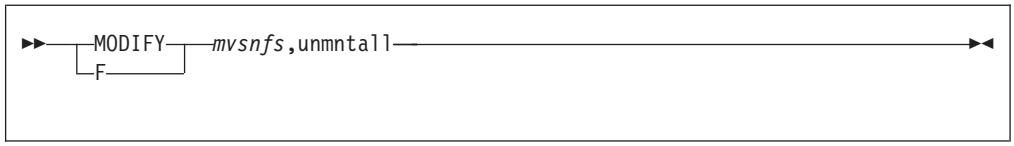

Unmnthfs operand

This operand causes the server to immediately unmount all z/OS UNIX mount points without stopping and restarting the NFS server. The general command syntax follows.

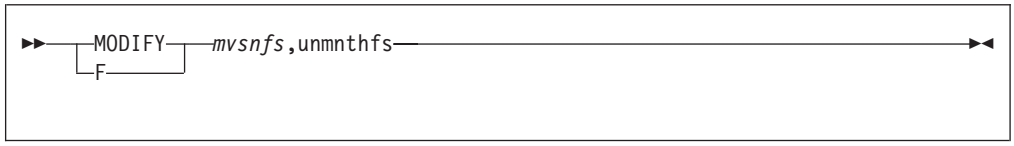

Unmntmvs operand

This operand causes the server to immediately unmount all MVS data set mount points without stopping and restarting the NFS server. The general command syntax follows.

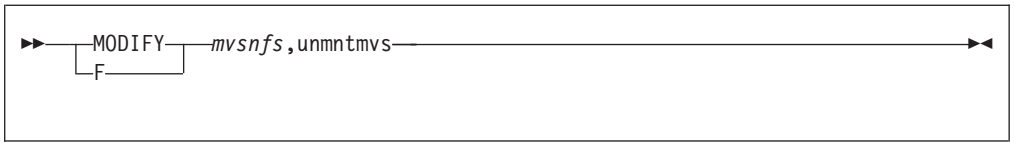

Entering operands of the modify command for diagnosis

These operands of the **modify** command can help you to collect data for diagnosing problems.

Note: All operands must be preceded with either:

```
MODIFY mvsnfs,
```

or

```
F mvsnfs,
```

where

mvsnfs Specifies the name of the procedure in your system that PROCLIB used to start up the server.

Each message is sent to the console, the component trace, or the data set that is pointed to by the NFSLOG1 or NFSLOG2 DD statements, or a combination of those destinations.

Flushctr operand

flushctr is an operator command that lets you flush the NFS component trace buffer records to an external writer. See "Using a z/OS component trace external writer" on page 233 for information on setting up an external writer for use with NFS component trace. Enter the operand as follows.

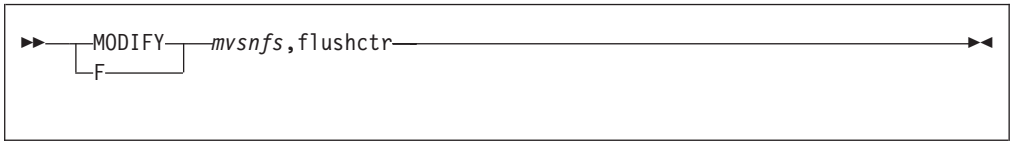

Flushlog operand

flushlog is an operator command that lets you flush the NFS message log to disk. This command enables a TSO/E user to browse all the log records that have been written by the NFS. Enter the operand as follows.

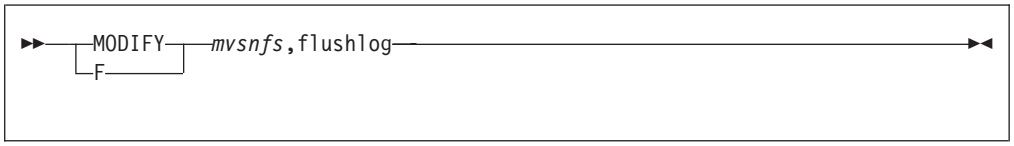

Listlock operand

This operand lists client NFS processes that hold locks for a specified file on the z/OS Network File System server. The list of lock holders can be used to diagnose problems with locking conflicts. The listlock operand lets you specify an MVS data

set, PDS or PDSE member, or z/OS UNIX file, and writes a message (GFSA791I or GFSA792I) for each unique userid and client host pair that holds locks for it. The general command syntax is as follows.

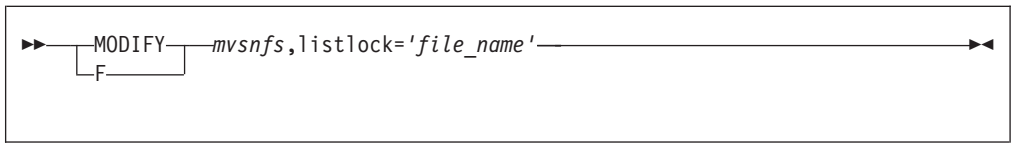

where

file_name
: Is the name of an MVS data set, z/OS UNIX file, PDS member, or PDSE member. If *file_name* references a PDS or a PDSE data set (does not include the member name), a syntax error will result. A z/OS UNIX file must be specified using the prefix identified in the NFS site attributes parameter HFS. For example, if the HFS(/hfs) site attribute was specified, LISTLOCK='/hfs/u/user' will indicate the z/OS UNIX file /u/user.

- If the specified file does not have any locks, message GFSA793I is issued to report that no locks exist.
- Since the lock information may only be reported to the server log data set or component trace buffer, and not back to the console, message GFSA794I is issued to indicate that the listlock function completed successfully.
- Before using this command, you might want to flush the log using the flushlog operand.
- Response messages are sent to the console if there are 10 or fewer locks to be reported. Otherwise, the response messages are only sent to the NFS log data set.

Log=msglevel operand

Use this operand to set the level of NFS logging messages to be collected. Enter the operand as follows.

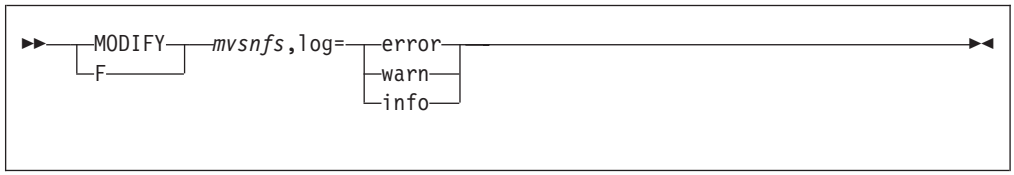

where

log=error
: Collects only error messages.

log=warn
: Collects error message and attention messages.

log=info
: Collects error, attention, and informational messages.

The diagnostic message level can also be set in the execution statement of the Network File System startup procedure.

For more information, see Appendix L, "Capturing diagnostic information using z/OS NFS log data sets and from other components," on page 379.

Smf operand

When the z/OS NFS server starts up, *smf* is set to off. You need to set *smf* to on if you want to collect the SMF records. The Network File System suspends writing SMF records when it detects an SMF error. The *smf* operand is used to resume or suspend the collection of SMF data. This operand has two parameters, ON and OFF. Enter the operand as follows.

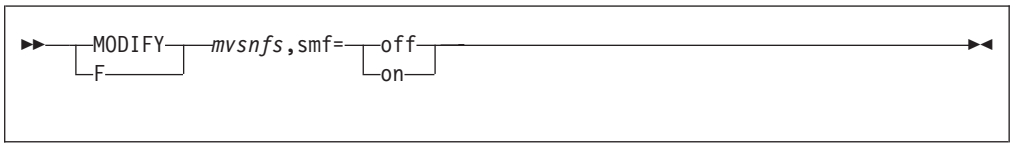

where

smf=on Resumes the collection of SMF records.

smf=off Suspends the collection of SMF records.

Version operand

To aid in diagnostics, the *version* operand identifies the precise release and maintenance levels of the z/OS NFS server modules, including the latest APAR installed. Enter the operand as follows.

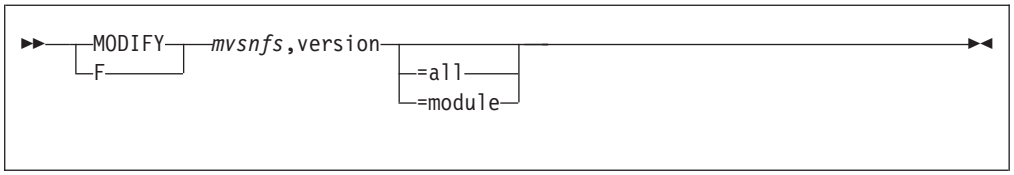

where

version Returns information about the release level of the z/OS NFS server and the latest APAR installed on the server.

version=all
Returns a list of all the modules in the z/OS NFS server with their current level information. At the end of the list is information about the latest APAR installed.

version=module
Returns information about the APAR level of the specified module.

Chapter 12. Installation-wide exits for the z/OS NFS server

This chapter helps you write installation-wide exits or replaceable modules that customize NFS processing. It contains product-sensitive programming interfaces that are provided by the server.

If there is no customization made to the installation-wide exit routines, we recommend that you remove them for better performance.

Requirements for NFS

Table 30 summarizes the installation-wide exit routines that are provided by the NFS.

Table 30. List of installation-wide exits

Source Module Name	Load Module Name	Parameter List Mapping Macro	Description
GFSAUXL	GFSAUXL	GFSAULOG GFSAUDSA	Login security exit
GFSAUXF	GFSAUXF	GFSAUSEC GFSAUDSA	File security exit

GFSAUDSA is a sample skeleton for a user storage block.

The above exits are shipped with the NFS and they contain the dummy skeleton code. The source modules reside in the prefix.NFSSAMP data set and the macros reside in prefix.NFSMAC data set. (The value of prefix depends on the installation.) Before modifying or replacing these exits, you should review the functions and processing of these exit routines carefully. These are basic requirements for all the NFS exit routines:

- Exit routines must reside in an authorized program library and be accessible by the z/OS LOAD macro. NFS and installation-wide exits receive control in problem state key 8. Installation-wide exits are run as an APF-authorized task, because Network File System is APF-authorized. As with any APF-authorized program, your exits should not be link-edited with APF-authorization. Only the main task, NFS, should have that link edit attribute.
- Exit routines must be link-edited with AMODE(31).

  The installation-wide exits are entered in AMODE(31) and can reside above or below the 16M line depending on the requirements of the installation-wide exits themselves.
- Exit routines must be reentrant.
- Exit routines must follow the standard z/OS register save and restore convention. The standard z/OS registers convention is:
  - Register 1 contains the address of the exit parameter list.
  - Register 13 contains the address of the caller's save area.
  - Register 14 contains the caller's return address.
  - Register 15 contains the address of the entry point for this exit routine. The server does not use return codes stored in register 15, but includes a parameter in the parameter list for exits that supply return codes.

© Copyright IBM Corp. 1991, 2006

Notes:

1. Address parameters have null value (0) if the related data does not exist.
2. The length of each field can be found in the macros shipped. Field length can be changed in the future.

Sample link-edit JCL

Use the sample JCL shown in Figure 25 to assemble and link-edit the GFSAUXL and GFSAUXF load modules:

```
//jobname  JOB (job_and_user_accounting_information)
//EXITASM  PROC M=
//ASM      EXEC PGM=ASMA90,
//              PARM='RENT'
//SYSPRINT DD   SYSOUT=*
//SYSIN    DD   DISP=SHR,DSN=source_library_name(&M)
//SYSLIB   DD   DISP=SHR,DSN=SYS1.MACLIB
//         DD   DISP=SHR,DSN=source_macro_library_name
//SYSLIN   DD   DISP=SHR,DSN=obj_library_name(&M)
//SYSUT1   DD   UNIT=SYSDA,
//              SPACE=(32000,(30,30))
//PEND
//         EXEC EXITASM,M=GFSAUXL
//         EXEC EXITASM,M=GFSAUXF
//stepname EXEC PGM=HEWL,
//         PARM='MAP,LIST,RENT,REUS'
//SYSPRINT DD SYSOUT=A
//SYSUT1   DD DSN=&SYSUT1,UNIT=SYSDA,SPACE=(CYL,(1,1))
//SYSLMOD  DD DSN=apf_library_name,DISP=OLD
//USER     DD DSN=obj_library_name,DISP=SHR
//SYSLIN   DD *
   INCLUDE USER(GFSAUXL)
   MODE    AMODE(31)
   ENTRY   GFSAUXL
   NAME    GFSAUXL(R)
   INCLUDE USER(GFSAUXF)
   MODE    AMODE(31)
   ENTRY   GFSAUXF
   NAME    GFSAUXF(R)
/*
```

Figure 25. Sample link-edit JCL for the NFS. This example uses High Level Assembler. To use Assembler H, replace `ASMA90` on the `EXEC` statement with `IEV90`. The rest of the JCL would be the same.

Storage blocks of the server exits

This section discusses how to use storage blocks of the NFS installation-wide exits.

Global exit block (GXB)

The GXB is obtained once during system initialization by the server login exit. The exit returns a word to the server. This word is referred to as 'the address of the GXB', but the system initialization exit might store any value in the word. The address of the GXB is returned to the server in the parameter list and passed back to the installation-wide exits in each subsequent call. This block contains user installation-wide exit data that is needed to communicate with the server.

The GXB can contain an area to save the user data needed for all sessions. The usage of this block is determined by the exit. Access to the Global Exit Block must be controlled by the user-written installation-wide exits to ensure that updates are serialized and do not interfere with each other. This block is shared with the file

security Installation-Wide Exit. The format of the GXB is entirely controlled by the login and file security installation-wide exits.

User exit block (UXB)

During a *Start of New User Session* request or a *User Login Request*, the exit can obtain a User Exit Block. The exit returns a word to the server. This word is referred to as 'the address of the UXB', but the system initialization exit might store any value in the word. The address of the UXB is returned to the server in the parameter list (depends on which request comes first), and is passed back to the installation-wide exits on each subsequent call related to this combination of machine and user IDs.

The UXB can contain an area to save the user data needed for this session. The usage of this block is determined by the exit. The exit is responsible for obtaining, and freeing access to these storage areas. This block is used by the login installation-wide exit and file security installation-wide exit. The format of the UXB is entirely controlled by the login and file security installation-wide exits.

Login installation-wide exit

The exit routine can invoke a customized authorization facility. The server mainline code can be set to perform Security Authorization Facility (SAF) checking, by specifying the **security** attribute in the attributes table.

If **security(saf)** or **security(safexp)** is specified in the attributes table and the exit routines exist, these exit routines get control first, and then SAF security checking gets control. If the exit routines fail the request, the entire request fails. If the exit routines process the request successfully, then the request is processed by the SAF checking. Similarly, if the SAF checking fails the request, the entire request fails.

If neither **security(saf)** nor **security(safexp)** is specified in the attributes table and the exit routine exists, this exit routine determines whether the request is successful or fails.

Figure 26 on page 176 shows the logic flow that determines which login checking routines are used.

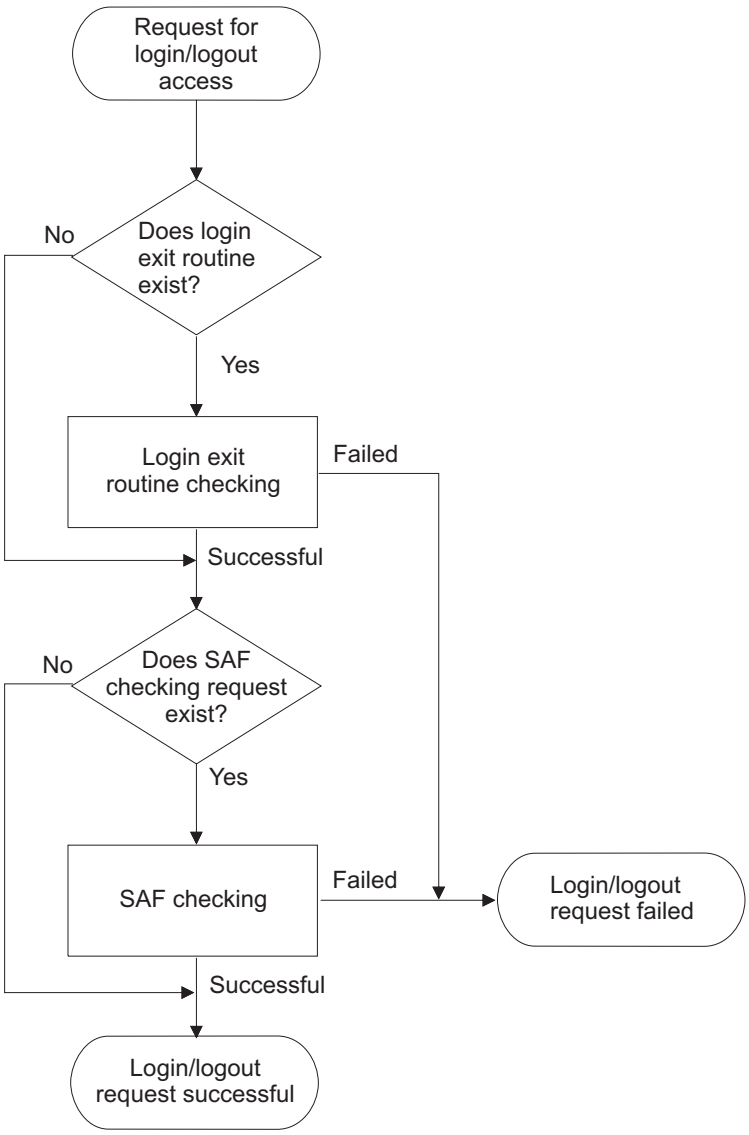

Figure 26. Determining which login checking routines are used

The login installation-wide exit has a parameter list which is passed from the server to the installation-wide exit. The login installation-wide exit can be invoked for any of the following conditions:

System initialization
Performed once during the initialization of the server and allows a Global Exit Block to be obtained. The GXB address is always returned to the installation-wide exits (see GXB in "Requirements of the login exit" on page 177). If this request fails, both the login installation-wide exit and the file security installation-wide exit are marked as non-existent.

New user session
Performed when a unique combination of UNIX UID and Internet address is detected. The exit might obtain a User Exit Block for use by later calls if the UXB does not exist.

Login
Performs user verification when a client tries to use either an **mvslogin**

command or a PCNFSD request. The exit might obtain a User Exit Block for use by later calls if the UXB does not exist.

Logout
Performs cleanup when a client tries to use the **mvslogout** command or a timeout occurs. Timeout is the value specified in the logout attribute in the attributes table. On a logout, the UXB is released. Logout can also be initiated by the login request.

System termination
Occurs once during the termination of the server, and causes the Global Exit Block to be freed.

Requirements of the login exit

Besides the requirements stated in "Requirements for NFS" on page 173, the login installation-wide exit routine must be link-edited with the name of GFSAUXL.

Options of the login exit

The login installation-wide exit routine can perform these functions:

- Obtain a Global Exit Block (GXB)
- Obtain a User Exit Block (UXB)

  If the User Exit Block already exists, its address is passed to the installation-wide exit routine; otherwise, 0 is passed. The installation-wide exit routine can accept or reject the request. If the request is accepted and the User Exit Block does not exist, you can allocate the User Exit Block and return the address in the LEDSXS parameter list.

Structure of the login exit message

The message supplied by the login exit is sent to the NFS server log data set and component trace, if activated. The message has the label name LEDSXD and consists of these two fields:

LEDSXCNT
1 byte message length, excluding the NULL character (X'00').

LEDSXMGS
81 byte NULL character string. The installation-wide exit can fill in this message field with the message ended by one or more bytes of X'00'.

If the format of the message is incorrect, both the login installation-wide exit and the file security installation-wide exit are marked as nonexistent, and a message (GFSA990I or GFSA991E) is sent to the z/OS NFS server log data set and to the console.

Contents of the login exit parameter list

The parameter list is mapped by macro and DSECT GFSAULOG. Table 31 on page 178 describes the contents of each field.

For configurations that use Internet Protocol Version 4, the client IP address is in field LEDSIA. If your configuration uses IPV6, refer to field LEDSIA6 for the client IP address. When an IPV6 address is provided in LEDSIA6, LEDSIA takes a value of -1.

If the **dhcp** site attribute is specified for the NFS server to use dynamic client IP addresses, the contents of the client IP address field will be correct when the exit

parameter list request is processed. However, the exit must not have any dependency on the persistence of the IP address value beyond the duration of the individual request. That IP address can change between requests.

Table 31. Format of login installation-wide exit routine parameter list

Field Name	Description	Contents
LEDSRQ	Request code	System request code set by the server before calling this installation-wide exit, for these conditions: **4** System initialization **8** System termination **12** Start of new user Session **16** User login request **20** Logout has been requested
LEDSRC	Return code	Codes generated and returned by the calls.
LEDSM	Client machine name	Character string ended by single byte containing X'00'.
LEDSIA	Client IP address	Number (32-bit Internet address). Contains -1 (X'0xFFFFFFFF) if an IPV6 address is provided in the LEDSIA6 field.'
LEDSU	UNIX Client user ID number	Number
LEDSG	UNIX Client group ID number	Number
LEDSMU	z/OS user ID	Character string padded with a blank at the end of the user ID,conforming to z/OS standards.
LEDSMG	z/OS group name	Character string padded with a blank at the end of the group name, conforming to z/OS standards.
LEDSXS	Address of UXB	Size and content are installation-dependent, generated at the start of a new user session or a user login request.
LEDSXG	Address of GXB	Size and content are installation-dependent, generated at system initialization.
LEDSXD	Message supplied by this exit routine	Message structure (see "Structure of the login exit message" on page 177).
LEDSVERS	z/OS version number	Number (*0x17* for z/OS V1R7)
LEDSIA6	Client IP address	Number (128-bit Internet address, for Internet Protocol V6 users)

Note: Character strings are in upper-case EBCDIC.

Login exit parameter list

This section describes how the Global Exit Block (GXB) or User Exit Block (UXB) is constructed and used between the login installation-wide exit and the NFS. A request code is set by the server before each call to this installation-wide exit routine. The installation-wide exit routine provides a return code for each event.

Request codes to the login exit

Table 32 shows the login installation-wide exit request codes and their meanings:

Table 32. Request codes to the login exit

Code	Meaning
4 (X'04')	System initialization
8 (X'08')	System termination

Table 32. Request codes to the login exit (continued)

Code	Meaning
12 (X'0C')	Start of new user session
16 (X'10')	User login request
20 (X'14')	User logout request

IBM might add new request codes in a future level of the server. To provide for this, consider making your exit set a return code 0 if it does not recognize the request code.

Return codes from the login exit

Table 33 shows the login installation-wide exit routine return codes and their meanings:

Table 33. Return codes from the login exit

Code	Meaning
0 (X'00')	Request successful
4 (X'04')	Request unsuccessful

System initialization routine of the login exit

The system initialization routine is called each time the server address space starts, and can acquire and initialize the GXB. Table 34 shows the codes and fields that are used for this routine:

Table 34. Codes and fields for system initialization

When Invoked	Field Name	Description	Contents	
On entry	LEDSRQ	Request code	**4**	
On exit	LEDSXD	Exit-supplied error message	Message Structure	
On exit	LEDSRC	Return code	**0**	Initialization successful
			4	Stop the NFS.
On exit	LEDSXG	Global Exit Block	Address	

Start of new user session routine of the login exit

The start of new user session routine is called when a new client machine-user combination is recognized by the server. The exit might acquire and initialize a UXB. Table 35 shows the codes and fields that are used for this routine:

Table 35. Codes and fields for start of new user session

When Invoked	Field Name	Description	Contents
On entry	LEDSRQ	Request code	**12**
On entry	LEDSM	Client machine name	EBCDIC character string
On entry	LEDSIA	Client IP address	Number
On entry	LEDSU	UNIX Client User ID number	Number
On entry	LEDSG	UNIX Client Group ID number	Number
On entry	LEDSXS	User Exit Block	0
On entry	LEDSXG	Global Exit Block	Value

Table 35. Codes and fields for start of new user session (continued)

When Invoked	Field Name	Description	Contents
On entry	LEDSIA6	Client IP address	Number (128-bit Internet address, for Internet Protocol V6 users)
On exit	LEDSXD	Exit-supplied error message	Message Structure
On exit	LEDSRC	Return code	**0** New session established **4** New session not established
On exit	LEDSXS	User Exit Block	Value, if return code in LEDSRC is 0

User login request routine of the login exit

This routine is called when the **mvslogin** command or PCNFSD request is used. The installation security system should be called to determine if the user is properly authorized. Table 36 shows the codes and fields:

Table 36. Codes and fields for user login request

When Invoked	Field Name	Description	Contents
On entry	LEDSRQ	Request code	**16**
On entry	LEDSM	Client machine name	Character string
On entry	LEDSIA	Client IP address	Number
On entry	LEDSU	UNIX Client User ID number	Number
On entry	LEDSG	UNIX Client Group ID number	Number
On entry	LEDSMU	z/OS User ID	Character string
On entry	LEDSMG	z/OS Group Name	Character string
On entry	LEDSXS	User Exit Block	Address or 0
On entry	LEDSXG	Global Exit Block	Address
On entry	LEDSIA6	Client IP address	Number (128-bit Internet address, for Internet Protocol V6 users)
On exit	LEDSXD	Exit-supplied error message	Message Structure
On exit	LEDSRC	Return code	**0** *Login* successful **4** *Login* failed
On exit	LEDSXS	User Exit Block	Value

User logout request routine of the login exit

This routine is used at logout to return the User Exit Block storage obtained at the start of the session and to perform any related logout processing. The codes and fields used are shown in Table 37.

Table 37. Codes and fields for logout request

When Invoked	Field Name	Description	Contents
On entry	LEDSRQ	Request code	**20**

Table 37. Codes and fields for logout request (continued)

When Invoked	Field Name	Description	Contents	
On entry	LEDSM	Client machine name	Character string	
On entry	LEDSIA	Client IP address	Number	
On entry	LEDSU	UNIX Client User ID number	Number	
On entry	LEDSG	UNIX Client Group ID number	Number	
On entry	LEDSMU	z/OS User ID	Character string	
On entry	LEDSMG	z/OS Group Name	Character string	
On entry	LEDSXS	User Exit Block	Address or 0	
On entry	LEDSXG	Global Exit Block	Address	
On entry	LEDSIA6	Client IP address	Number (128-bit Internet address, for Internet Protocol V6 users)	
On exit	LEDSXD	Exit-supplied error message	Message Structure	
On exit	LEDSRC	Return code	**0**	*Logout* successful
On exit	LEDSXS	User Exit Block	**0**	UXB is released

System termination routine of the login exit

This routine is used at server termination to release the GXB storage. All users are automatically logged off. The codes and fields used are shown in Table 38.

Table 38. Codes and fields for system termination

When Invoked	Field Name	Description	Contents	
On entry	LEDSRQ	Request code	**8**	
On entry	LEDSXG	Global Exit Block	Address	
On exit	LEDSXD	Exit-supplied error message	Message Structure	
On exit	LEDSRC	Return code	**0**	Exit termination successful
On exit	LEDSXG	Global Exit Block	**0**	GXB is released

File security installation-wide exit

The file security installation-wide exit routine verifies that a user is authorized to access a data set or data set member with the access mode requested. If the request from allocation, write, read, or access does not have permissions set up, then the exit routine gets control.

This exit applies only to MVS data set access, not to z/OS UNIX file access.

The permissions set up by the file security exit can be overridden by the SAF checking. If the exits allow access and there is no SAF checking, the permissions

remain in effect until logout. The server does not call again for the same access before logout. The server gets the access mode or permissions before any of the other three types of calls.

If **security(saf)** or **security(safexp)** is specified in the attributes table and the exit routine exists, this exit routine gets control first, and then SAF security checking gets control. If the exit routines fail the request, the entire request fails. If the exit routines process the request successfully, then the request is processed by the SAF checking. Similarly, if the SAF checking fails the request, the entire request fails. If the SAF checking is successful, the file permissions from the SAF checking are set up for the request.

If neither **security(saf)** nor **security(safexp)** is specified in the attributes table and the exit routine exists, this exit routine determines the permissions.

Figure 27 on page 183 shows the logic flow determining which file security checking routines are used.

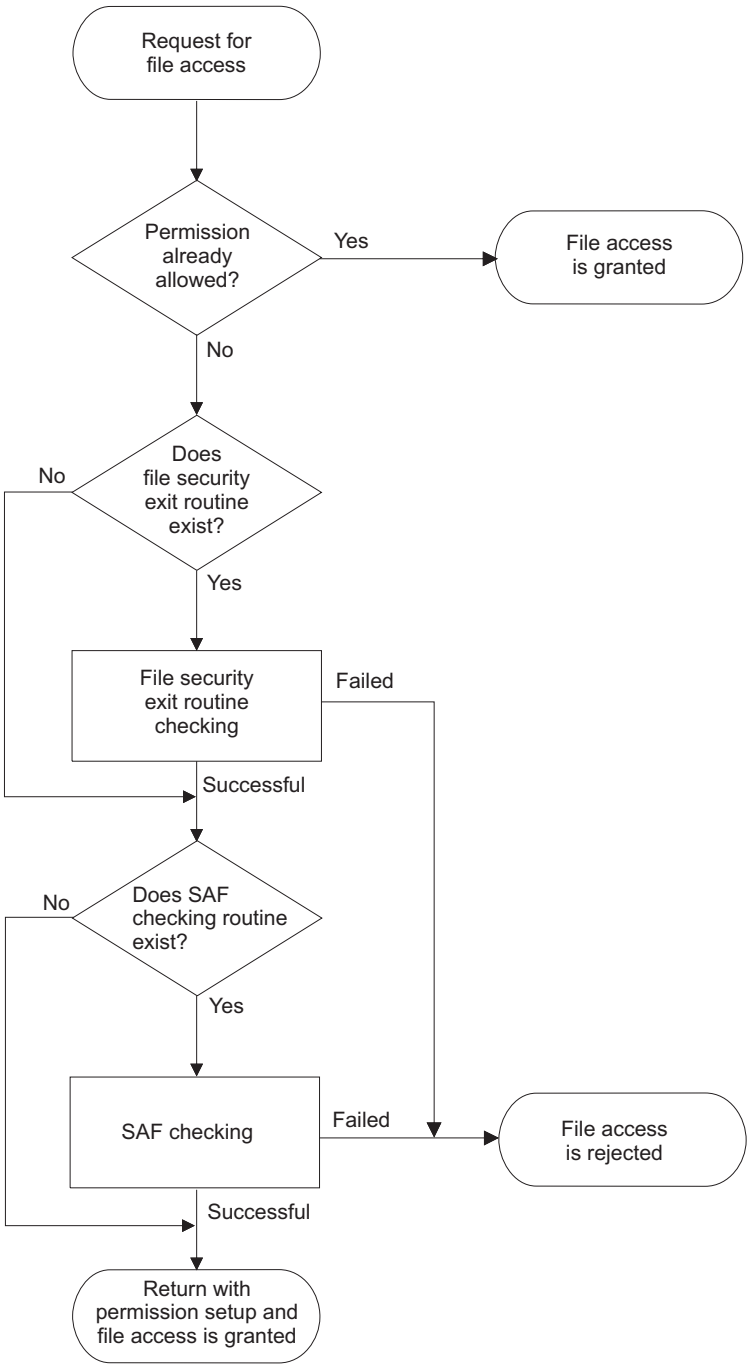

Figure 27. Determining which file security checking routines are used

The file security installation-wide exit uses the parameter list which is used by the server installation-wide exit. This exit is called for at these times.

Data Set Allocation
When a client user creates, renames, or deletes a data set.

Data Set Write
When a client user first writes to an MVS data set or data set member.

Data Set Read
When a logged in user tries to first read from an MVS data set or data set member.

Getting Access Mode or Permissions

This function is called first to set up permission for the read, write, create, delete, or rename data set request. The server needs to obtain the access mode or the permissions that a user has for a data set. This function is only called when SAF is not being used.

The following show access modes and what they permit.

Allocate	Read, write, create, delete, or rename the data set.
Write	Read or write the data set
Read	Read the data set

A return code is set by the installation-wide exit indicating whether the request is allowed. The file security installation-wide exit is not called at server startup or shutdown.

Requirements of the file security exit

Besides "Requirements for NFS" on page 173, the following requirements must be fulfilled.

- The file security installation exit must be link-edited with the name of GFSAUXF.
- The login exit must exist (GFSAUXL).
- If the GXB or UXB are to be used, the login exit must obtain them.
- The UXB and GXB are shared with the login installation-wide exit.

Structure of the file security exit message

The message supplied by the file security exit is sent to the NFS server log data set and component trace, if activated. The message has the label name FEDSXD and consists of the following two fields.

FEDSXCNT
: 1 byte message length, excluding the NULL character (X'00').

FEDSXMGS
: 81 byte NULL character string. The installation-wide exit can fill in this message field with an error message ended by one or more bytes of X'00'.

If the format of the message is incorrect, the file security installation-wide exit is marked as nonexistent, and a message (GFSA990I or GFSA991E) is sent to the NFS server log data set and to the console, and to component trace if activated.

Contents of the file security exit parameter list

The file security installation-wide exit parameter list is mapped by macro and DSECT GFSAUSEC. Table 39 on page 185 describes each field.

For configurations that use Internet Protocol Version 4, the client IP address is in field FEDSIA. If your configuration uses IP version 6, refer to field FEDSIA6 for the client IP address. When an IPv6 address is provided in FEDSIA6, FEDSIA takes a value of -1.

If the **dhcp** site attribute is specified for the NFS server to use dynamic client IP addresses, the contents of the client IP address field will be correct when the exit parameter list request is processed. However, the exit must not have any dependency on the persistence of the IP address value beyond the duration of the

individual request. That IP address can change between requests.

Table 39. Format of the parameter list for the file security installation-wide exit

Field Name	Description	Contents
FEDSRQ	Request code	System request code set by the server before each call to this exit for these conditions. **4** Validate allocate request **8** Validate write request **12** Validate read request **16** Return the current settings of the security permissions
FEDSRC	Return code	Codes are generated and returned by the exit routine.
FEDSM	Client system name	Character string ended by a single byte containing X'00'.
FEDSIA	Client IP address	Number (32-bit Internet address). Contains -1 (X'0xFFFFFFFF) if an IPV6 address is provided in the FEDSIA6 field.'
FEDSU	Client user ID number	Number
FEDSG	Client group ID number	Number
FEDSMU	z/OS user ID	Character string padded with a blank at the end of the user ID, conforming to z/OS standards.
FEDSDN	MVS data set name	Character string ended by a single byte containing X'00'. This conforms to z/OS standards.
FEDSMN	MVS data set member name	Character string ended by a single byte containing X'00'. This conforms to z/OS standards.
FEDSXS	Address of UXB	Size and contents are installation dependent. Generated at the start of a user session or user login by the *login* exit routine.
FEDSXG	Address of GXB	Size and content are installation dependent. Generated at server initialization by the *login* exit routine.
FEDSTYPE	File Type	Address of fullword integer defining the data set organization: **0** Not used **1** Sequential **2** Partitioned **3** Direct **4** Reserved **5** VSAM (if unable to classify further) **6** VSAM ESDS **7** VSAM RRDS **8** VSAM KSDS **9** Reserved **10** Reserved
FEDSRPM	UNIX Permissions	UNIX file modes of type mode_t as defined by POSIX.
FEDSXD	Message supplied by this exit routine	Address of the message structure (see "Structure of the file security exit message" on page 184).
FEDSVERS	z/OS version number	Number (*0x17* for z/OS V1R7)

Table 39. Format of the parameter list for the file security installation-wide exit (continued)

Field Name	Description	Contents
FEDSIA6	Client IP address	Number (128-bit Internet address, for Internet Protocol version 6 users)

Note: Character strings are in upper-case EBCDIC.

File security exit parameter list

These sections each describe how the data in the parameter list is used by the server and the file security installation-wide exit routine. A request code is set by the server before each call to this installation-wide exit routine.

Each topic below describes an event, for which some fields are set on entry. The file security installation-wide exit provides a return code for each event.

Request codes to the file security exit

Table 40 shows the request codes to the file security installation-wide exit:

Table 40. Request codes to the file security exit

Code	Meaning
4 (X'04')	Validate allocate request
8 (X'08')	Validate write request
12 (X'0C')	Validate read request
16 (X'10')	Return the current settings of the security permissions

Return codes from the file security exit

Table 41 shows the request codes returned by the file security installation-wide exit:

Table 41. Return codes from the file security exit

Code	Meaning
0 (X'00')	Request successful
4 (X'04')	Request unsuccessful

Validate allocate request routine of the file security exit

This routine is called when a user allocates, renames, or deletes a data set or data set member. The codes and fields used are shown in Table 42.

Table 42. Codes and fields for validate allocate request

When Invoked	Field Name	Description	Contents
On entry	FEDSRQ	Request code	**4**
On entry	FEDSM	Client machine name	Character string
On entry	FEDSIA	Client IP address	Number
On entry	FEDSU	Client user ID number	Number
On entry	FEDSG	Client group ID number	Number
On entry	FEDSMU	z/OS user ID	Character string
On entry	FEDSDN	MVS data set name	Character string
On entry	FEDSMN	MVS data set member name	Character string

Table 42. Codes and fields for validate allocate request (continued)

When Invoked	Field Name	Description	Contents	
On entry	FEDSXS	User Exit Block	Address	
On entry	FEDSXG	Global Exit Block	Address	
On entry	FEDSTYPE	File type	Number	
On entry	FEDSIA6	Client IP address	Number (128-bit Internet address, for Internet Protocol version 6 users)	
On exit	FEDSXD	Exit-supplied error message	Message Structure	
On exit	FEDSRC	Return code	**0**	Access allowed
			4	Access denied

Validate write request routine of the file security exit

This routine is called when a user writes to an MVS data set or data set member. The codes and fields used are shown in Table 43.

Table 43. Codes and fields for validate write request

When Invoked	Field Name	Description	Contents	
On entry	FEDSRQ	Request code	**8**	
On entry	FEDSM	Client machine name	Character string	
On entry	FEDSIA	Client IP address	Number	
On entry	FEDSU	Client User ID number	Number	
On entry	FEDSG	Client Group ID number	Number	
On entry	FEDSMU	z/OS user ID	Character string	
On entry	FEDSDN	MVS data set name	Character string	
On entry	FEDSMN	MVS data set member name	Character string	
On entry	FEDSXS	User Exit Block	Address	
On entry	FEDSXG	Global Exit Block	Address	
On entry	FEDSTYPE	File type	Number	
On entry	FEDSIA6	Client IP address	Number (128-bit Internet address, for Internet Protocol version 6 users)	
On exit	FEDSXD	Exit-supplied error message	Message Structure	
On exit	FEDSRC	Return code	**0**	Access allowed
			4	Access denied

Validate read request routine of the file security exit

This routine is called when a user reads from an MVS data set or data set member. The codes and fields used are shown in Table 44.

Table 44. Codes and fields for validate read request

When Invoked	Field Name	Description	Contents
On entry	FEDSRQ	Request code	**12**
On entry	FEDSM	Client machine name	Character string

Table 44. Codes and fields for validate read request (continued)

When Invoked	Field Name	Description	Contents
On entry	FEDSIA	Client IP address	Number
On entry	FEDSU	Client user ID number	Number
On entry	FEDSG	Client group ID number	Number
On entry	FEDSMU	z/OS user ID	Character string
On entry	FEDSDN	MVS data set name	Character string
On entry	FEDSMN	MVS data set member name	Character string
On entry	FEDSXS	User Exit Block	Address
On entry	FEDSXG	Global Exit Block	Address
On entry	FEDSTYPE	File type	Number
On entry	FEDSIA6	Client IP address	Number (128-bit Internet address, for Internet Protocol version 6 users)
On exit	FEDSXD	Exit-supplied error message	Message Structure
On exit	FEDSRC	Return code	**0** Access allowed **4** Access denied

Return security permissions routine of the file security exit

This routine is called to query a data set access mode or permission. The codes and fields used are shown in Table 45.

Table 45. Codes and fields for return security permissions

When Invoked	Field Name	Description	Contents
On entry	FEDSRQ	Request code	**16**
On entry	FEDSM	Client machine name	Character string
On entry	FEDSIA	Client IP address	Number
On entry	FEDSU	Client user ID number	Number
On entry	FEDSG	Client group ID number	Number
On entry	FEDSMU	z/OS user ID	Character string
On entry	FEDSDN	MVS data set name	Character string
On entry	FEDSMN	MVS data set member name	Character string
On entry	FEDSXS	User Exit Block	Address
On entry	FEDSXG	Global Exit Block	Address
On entry	FEDSTYPE	File type	Number
On entry	FEDSIA6	Client IP address	Number (128-bit Internet address, for Internet Protocol version 6 users)
On exit	FEDSXD	Exit-supplied error message	Message Structure
On exit	FEDSRC	Return code	**0** Exit processing successful **4** Exit failed

Table 45. Codes and fields for return security permissions (continued)

When Invoked	Field Name	Description	Contents
On exit	FEDSRPM	New permissions in mode_t format	UNIX file modes of type mode_t as defined by POSIX

Part 4. Performance Tuning

Chapter 13. Performance tuning in the NFS environment 193
What is performance tuning? 193
How is performance characterized? 193
What is the NFS environment? 194
How to tune for performance 194
Impact of the NFS protocol on performance . . . 196

Chapter 14. Optimizing the NFS environment 199
Network performance tuning 199
NFS client system performance tuning 200
NFS server system performance tuning. 201
z/OS constraints 202

Chapter 15. Evaluating z/OS NFS performance 205
Evaluating throughput 205
Single process throughput 205
Multiple process throughput 206
Multiple client throughput 206
Evaluating NFS command response time 206
Evaluating CPU utilization 206
Collecting server usage data 207

Chapter 16. Tuning the z/OS NFS server . . . 209
Data set creation attributes 209
Block size and record length 209
Record format 210
Data set organization and data set type. . . . 210
Processing attributes 211
Character translation 211
File size determination 211
Data set timeout specification 213
Accessing migrated files. 213
Asynchronous z/OS UNIX processing 214
Site attributes 214
Buffer usage and caching 214
Ordering out-of-sequence data. 215
Storage considerations 217
Subtasking 217

Chapter 17. Tuning the z/OS NFS client. . . . 219
Caching 219
Dynamicsizeadj. 219
Bufhigh 220
Biod 220
Readahead 220
Delaywrite 220
Vers 220
Wsize and rsize 221

© Copyright IBM Corp. 1991, 2006

Chapter 13. Performance tuning in the NFS environment

This chapter explains performance tuning within the context of the NFS client-server environment. It provides guidelines on the performance expectation of the underlying hardware and software products on which an NFS client-server implementation is dependent. These guidelines specifically note the limitations to performance tuning and describe processes to tune z/OS NFS.

What is performance tuning?

In general, performance tuning improves the price to performance ratio for a system or set of services by reallocating the available computing, network, or storage resources. The reallocation of these resources not only improves the performance for a particular load of work, but also accommodates an increase in the amount of work to be performed with minimal acquisition of additional resources. The information acquired from performance tuning can also be an important basis for long range capacity planning.

The following might be some possible reasons to do performance tuning.

- Access more data over existing networks
- Improve response time for particular applications or groups of users
- Better utilize the available storage capacity
- Minimize cost for additional services or functional capability

How is performance characterized?

Performance is the manner in which a process, system, processor, network, or device behaves for a particular load or unit of work. To measure, or quantify, performance, we monitor the length of time for a unit of work to complete. If units of work are being shared, we monitor the amount of time waiting for a resource to be available to perform a unit of work. A unit of work is a specific activity or action that we expect a process, system, processor, network, or device to perform. This could be something as granular as an I/O request, sending or receiving a buffer of data over a network, or processing an NFS request. We are frequently interested in the performance of a particular set of work activities which we will refer to as a load of work, or workload.

When a particular workload has been identified for performance measurement, we can determine the performance metrics, or units of measurement, that are relevant to that workload. Some examples of the following performance metrics that might be used in reference to performance tuning for z/OS NFS are:

throughput	Number of units of work completed per unit of time.
response time	Elapsed time necessary to complete a unit of work.
CPU time	Amount of time the processor spends executing instructions.
instructions per byte	Total number of CPU instructions to process data divided by the number of bytes of data processed.
aggregate throughput	Sum of the throughput measurements for multiple processes.
NFS operations per second (NFSops)	Number of NFS requests processed by an NFS server divided by the total time in seconds to process the requests.

© Copyright IBM Corp. 1991, 2006

Measuring performance metrics can be as simple as using your watch to time the execution of a particular command or as complex as using specialized hardware and software tools to monitor and extract a diversity of performance metrics. Chapter 15, "Evaluating z/OS NFS performance," on page 205 will address some methods that can be used to evaluate the performance of z/OS NFS. The method selected will depend on the complexity of the workload and the monitoring tools available at your installation.

What is the NFS environment?

The NFS environment includes the NFS client system(s), the mix of networks available, the NFS server system(s), and the manner in which they are configured. While this guide is intended for z/OS NFS, it is important to know the performance limitations within the NFS environment to determine the necessity of tuning z/OS NFS.

NFS client systems range from single user desktop personal computers to large scale processors with many users. These NFS client systems typically support multiple applications as well. Clearly, the NFS client system resources will be shared between its users and/or applications. These resources include available physical storage, memory, processing capability, and network access. The NFS client is one application, with a possibility for many users, that must share the memory, processing, and network resources in exchange for providing access to additional physical storage on other systems. The degree to which the NFS client application must share such resources will affect performance for NFS client users and applications.

The NFS server system, like the NFS client system, must share resources with other users and applications. NFS server application performance will be affected by the amount of contention over system resources and by the priority established for the NFS server application. The overall performance of an NFS server is also influenced by the number of NFS clients for which it provides services.

The network(s) over which NFS clients access NFS servers also affect overall performance in the NFS environment. Such networks can be homogenous, consisting of a single network medium, or heterogeneous, consisting of a mix of network mediums. Each network medium type has an expected maximum capacity, or *bandwidth*. For instance, the capacity of a Fast Ethernet Ring network may be 16 megabits per second (Mbps) or 4 Mbps, and the capacity of a Fiber Distributed Data Interface (FDDI) network is 100 Mbps. When different network mediums are combined in a more complex network environment, the capacity for a fixed route over the network is limited by the network segment with the smallest capacity. For example, a route over a network consisting of both 4 and 16 Mbps Token Rings and a FDDI backbone will have a maximum capacity equivalent to that of the Token Ring, or 4 Mbps. When bridges, routers, and gateways are included in a network configuration, their capacity must also be considered. Such devices must also be considered when tuning performance in a network environment, particularly if a device does not support increased network buffer sizes.

How to tune for performance

Given the complexity of the NFS environment, it is important to establish a methodology for tuning performance. The following steps provide a guideline that highlights particular areas relevant to the NFS environment. Implementing the guideline may involve more than one person or support organization.

Identify Performance Requirements: Before you begin performance tuning in general, and particularly z/OS NFS, determine those areas where performance is unsatisfactory. This is a good time to establish more precise performance requirements. As users and application requirements are identified, it can be advantageous to rank or group them according to their requirements.

Know the NFS Environment: In the previous section, the NFS environment was discussed. It is very important to fully understand the performance of the existing NFS environment, particularly that of the network. Such analysis can eliminate unnecessary tuning of the NFS client and server systems.

Establish Performance Objectives: Once the performance requirements and the NFS environment are known, you are in a position to define and prioritize your performance objectives. The performance objectives should be specified in a manner that is quantifiable and measurable. Keep in mind that you will need an executable workload or test scenario to evaluate the effectiveness of your performance tuning.

Define Workloads and Test Scenarios: You may already have workloads or test scenarios depending on how requirements and/or objectives were defined. However, these may be unwieldy or impractical to use for your performance tuning purposes. Therefore, it is advantageous to spend some time initially to define some simplified test cases that can be executed in a repeatable fashion and with as much control as is feasible. Simple test cases are also useful in diagnosing performance problems, particularly to locate an area within the NFS environment that may be impacting performance.

Select Monitoring Tools: While you may only be interested in the performance of z/OS NFS, you will find it useful to have access and familiarity with a variety of performance hardware and software monitoring tools. Not only will such knowledge assist you in collecting data to evaluate performance, but it will also help you to identify areas that may be impacting the performance of z/OS NFS. Minimally, a set of monitoring tools must be identified to collect the data upon which performance tuning decisions will be made and to determine the effectiveness of tuning.

Collect Performance Data: At this point you should be ready to begin collecting performance data. Initial measurements will be the starting point, or *baseline*, that will be used to evaluate the effect of your tuning. Since there can be a significant variation in network and system performance, it is prudent to repeat a performance measurement to establish the degree of variation inherent in the measurement. Doing this will provide a sense of whether or not future tuning is really affecting performance or simply normal variation.

While it may be convenient to collect performance data when systems and networks are idle, this is probably not practical. However, it is useful to collect data during peak and low activity periods. If user or application requirements are related to a specific time period, or sensitive to other system and network activity, data should be collected for these periods of time, as well. Remember that redistributing the workload may be the most cost effective approach to performance tuning.

Evaluate Performance Data: This may seem like an obvious step. However, keep in mind that the NFS environment may be quite complex. The more complex the NFS environment and your test cases are, the more data there will be to evaluate. You may also have both client and server data to evaluate as well as network data.

As you begin evaluating performance data you have collected, look for areas or opportunities where performance could be improved. Before attempting to tune z/OS NFS, you should determine if performance is primarily impacted by the NFS server system, the NFS client system, or the network itself. In fact, you may determine that additional data must be collected prior to any performance tuning. You may also discover evidence that configurations and parameter settings within the evaluation environment are not optimal.

The output of this step is a list of changes that you believe will positively affect performance. As you make this list, identify the impact or cost associated with a change. It is also useful to identify any resources that are heavily utilized or that have contending requirements. This will help you to prioritize and ultimately select the changes you will make or recommend.

Tune Your NFS Environment: Make one change from your list of possible tuning changes at a time. Measure and evaluate the effect of that change before making any other changes. Pay particular attention to any impact on heavily utilized resources that may have been previously identified in addition to newly exposed resources that now appear to be impacted. Since performance tuning typically involves a trade off in resource utilization, make sure you have not inadvertently caused a performance problem elsewhere. Also, before deciding to implement a change, consider whether any observed changes are due to the tuning change you've made or simply a result of normal measurement variation. Repeat this step as necessary.

Impact of the NFS protocol on performance

Command response time is of particular importance to NFS client users. The longer a user waits for the results of a particular command, the more important this will become. The nature of transparent access with the NFS protocol results in users not necessarily being aware of the impact on performance caused by the network and the NFS server system.

Also, while users are generally knowledgeable of the commands supported by the NFS client operating system, they may have no knowledge of the NFS commands, or procedures, that are executed as a result of one simple user command. In fact, one user command typically results in execution of multiple NFS commands.

Another impact on command response time is the NFS protocol itself; version 2 and version 3 of the NFS protocol are intended to be as stateless as possible. This means that a stateless server operates correctly without maintaining any protocol state information for its clients. A stateless protocol was originally selected to minimize the probability of data losses due to a server crash. In NFS version 4, some state information is introduced into the protocol (for example, the open/close, lock, and setclientid operations).

For NFS version 2 and version 3 protocols, the stateless nature of the NFS server places the responsibility of keeping track of NFS commands on the NFS client. To do this NFS client implementations generally wait a period of time for a response to a particular NFS command. If a response is not received within this period of time, the NFS client will retransmit the NFS command. This process is repeated for a fixed number of times or until a response is received.

NFS servers and clients have typical methods of queuing requests and responses as part of the underlying protocol layers. When these queues are full, incoming requests or responses are dropped. The NFS client and server do not know when

responses or requests are dropped. Both rely on the stateless protocol whereby the client will eventually retransmit the NFS request again. Under these conditions the NFS client is waiting for a response that will never be received. Clearly, waiting to retransmit an NFS command, particularly multiple times, will negatively affect the response time for the initial user command.

To determine whether or not NFS client users are being impacted by the situation previously described, most NFS client implementations provide a **nfsstat** command to monitor NFS statistics. Figure 28 shows the output from the ***nfsstat -c*** command.

```
USER1:/u/user1:>nfsstat -c

Client rpc:
Connection oriented
calls      badcalls   badxids    timeouts   newcreds   badverfs   timers
0          0          0          0          0          0          0
nomem      cantconn   interrupts
0          0          0
Connectionless
calls      badcalls   retrans    badxids    timeouts   newcreds   badverfs
0          0          0          0          0          0          0
timers     nomem      cantsend
0          0          0

Client nfs:
calls      badcalls   clgets     cltoomany
0          0          0          0
Version 2: (0 calls)
null       getattr    setattr    root       lookup     readlink   read
0 0%       0 0%       0 0%       0 0%       0 0%       0 0%       0 0%
wrcache    write      create     remove     rename     link       symlink
0 0%       0 0%       0 0%       0 0%       0 0%       0 0%       0 0%
mkdir      rmdir      readdir    statfs
0 0%       0 0%       0 0%       0 0%
Version 3: (0 calls)
null       getattr    setattr    lookup     access     readlink   read
0 0%       0 0%       0 0%       0 0%       0 0%       0 0%       0 0%
write      create     mkdir      symlink    mknod      remove     rmdir
0 0%       0 0%       0 0%       0 0%       0 0%       0 0%       0 0%
rename     link       readdir    readdir+   fsstat     fsinfo     pathconf
0 0%       0 0%       0 0%       0 0%       0 0%       0 0%       0 0%
commit
0 0%
Version 4: (0 calls)
null       getattr    setattr    lookup     access     readlink   read
0 0%       0 0%       0 0%       0 0%       0 0%       0 0%       0 0%
write      create     mkdir      symlink    mknod      remove     rmdir
0 0%       0 0%       0 0%       0 0%       0 0%       0 0%       0 0%
rename     link       readdir    statfs     finfo      commit     open
0 0%       0 0%       0 0%       0 0%       0 0%       0 0%       0 0%
confirm    downgrade  close      lock       locku      lockt      setclid
0 0%       0 0%       0 0%       0 0%       0 0%       0 0%       0 0%
renew      clid_cfm   secinfo    release_lo replicate
0 0%       0 0%       0 0%       0 0%       0 0%
```

Figure 28. Displaying NFS client rpc and NFS statistical information

When using the **nfsstat** command, users should be aware that results are cumulative. These statistics may be reset with the -*z* option of the nfsstat command. Also, the nfsstat-*c* command provides statistics for all NFS client activity, which makes it possible to access files on more than one NFS server. When using this command to query the NFS client statistics for z/OS NFS, make sure that all NFS client access is for z/OS NFS only.

The nfsstat-*c* command provides two types of statistics, *client rpc* and *client nfs*. The *client rpc* statistics provide an indication of NFS performance from the perspective

of that particular NFS client. In general, if the *timeout* value is more than 0.2% of the total number of *rpc* calls, then some performance tuning is necessary.

If the *timeout* value is essentially equivalent to the *retrans* value, the NFS client is waiting on the NFS server. While you can increase the *timeout* option of the **mount** command to reduce the number of retransmissions, this will not improve the perceived responsiveness of the NFS server.

The *badcall* and *badxid* statistics are generated when information is lost or dropped somewhere between the NFS client and the NFS server. This can happen when processes are interrupted and are not necessarily indicative of a performance problem unless they are disproportionately high or persist. On the NFS client system the **netstat**-*s* command provides additional statistics that may be useful in tuning at the network level. For the z/OS NFS, the z/OS TCP/IP netstat command provides information that may be useful for tuning as well.

The client NFS statistics provide a distribution of the NFS procedure calls made by users and applications on that particular NFS client. For a typical NFS client you will probably see a larger percentage of *getattr* and *lookup* calls. These calls are made whenever an NFS file is initially accessed. Directory listing is a common user activity that will generate these NFS calls.

The write and read statistics can also provide insight into the manner of NFS access on an NFS client, such as read or write biases. If the percentage of read and write NFS calls is higher than the percentage of getattr and lookup NFS calls, the client is probably accessing relatively large files as opposed to accessing many smaller files. Conversely, if the percentage of getattr and lookup NFS calls is higher than the percentage of read and write NFS calls, you might want to investigate whether users are querying directories or accessing relatively small files. In the latter case you might further determine if such files should be stored on an NFS server or the NFS client system.

Chapter 14. Optimizing the NFS environment

This chapter explores areas that may impact the overall performance within an NFS environment. The focus will be on various network and NFS client and server system parameters which may affect NFS performance. Specific product documentation should be consulted for additional detail.

Network performance tuning

Table 46 contains network performance tuning information.

Table 46. Network performance tuning symptom and action information

Symptom	Action
Data transfer rates significantly less than network capacity	**ACTION 1:** *Know your network topology* - "What is the NFS environment?" on page 194 briefly introduces the role of the network in the NFS environment. Clearly, transferring data over congested networks results in poor performance. Therefore, when a performance problem is encountered in the NFS environment, it is useful to determine whether or not the network is the source of the problem prior to investigating other alternatives. While there are products specifically designed to monitor and report on network activity, knowing your particular network topology may be sufficient for initial tuning. Figure 29 shows a simple network topology with two 100 Mbps Fast Ethernet Rings connected to a Gbps Fast Ethernet backbone network. In this example, access from NFS client A, on a 100 Mbps Fast Ethernet Ring, to NFS server B, directly connected to the Gbps Fast Ethernet backbone, would be as limited by network bandwidth as access to NFS server A. On the other hand, access from NFS client B, directly connected to the Gbps Fast Ethernet backbone, would be limited by the load on the Gbps Fast Ethernet network and the capacity of NFS server B.

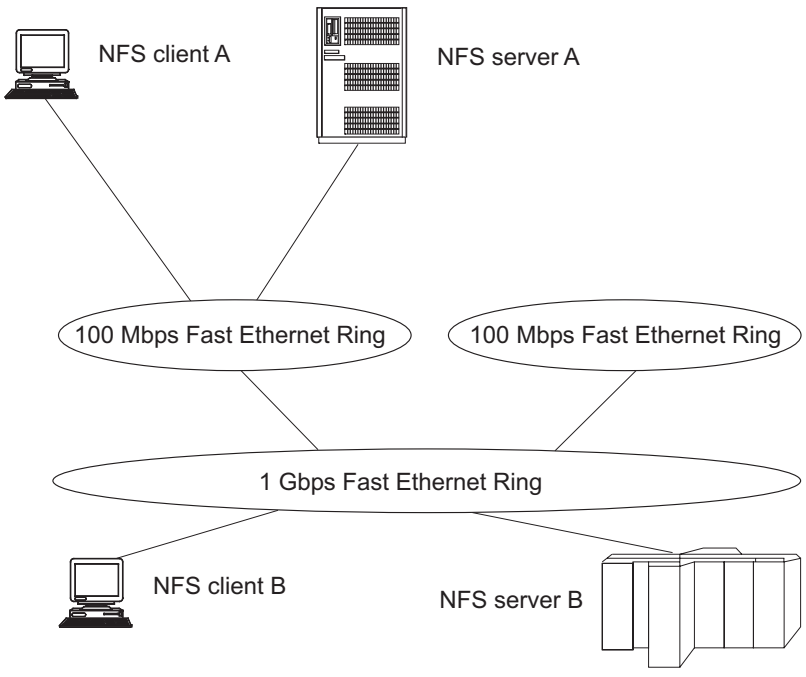

Figure 29. Sample network topology

© Copyright IBM Corp. 1991, 2006

There are other factors to consider in addition to the theoretical capacity of a network configuration. For instance, there are differences in the capability of network adaptors and network controllers. Overall network performance can also be influenced by the devices connecting network segments or subnets, particularly in terms of packet size limits and UDP checksum processing. While it may initially appear that there is no network constraint, there may be a device, adaptor, or network segment that is not performing at expected levels.

With this type of investigation, you may already have discovered that your primary constraint is a congested network. If so, you may elect to pursue such alternatives as:

- Increasing network bandwidth
- Changing network topology
- Scheduling applications during low system usage
- Modifying applications to reduce data transfer

ACTION 2: *Monitor network activity* - It may be necessary to monitor network activity over a period of time to determine what is causing a performance problem. In addition to whatever monitoring your network administrator may be able to provide, the **netstat -s** command, available on many NFS client platforms, may provide some insight as well. This command reports such statistics as bad checksums, dropped fragments, non-forwardable packets, various timeouts, and socket buffer overflows. In some cases, it may be necessary to use a network analyzer to determine exactly what is being sent over a network segment. While it can be difficult to analyze congested networks, network analyzers occasionally capture evidence that may indicate a problem with the NFS client or server system.

NFS client system performance tuning

Table 47 contains NFS client system performance tuning symptom and action information.

Table 47. NFS client system performance tuning symptom and action information

Symptom	Action
Excessive NFS retransmissions	A high number of NFS retransmissions can be detected with the **nfsstat -c** command. See section “Impact of the NFS protocol on performance” on page 196 for more information. A general rule of thumb is that the number of retransmissions should not be more than 2% of the total calls. Use **traceroute** *hostname* command to trace the route that IP packets take to a network host. This helps isolate where the bottleneck may occur.

Symptom	Action
	ACTION 1: *Increase network queue length* - If the **nfsstat -c** command shows evidence of bad calls in addition to timeouts, or the **nfsstat -s** indicates that packets or fragments are being dropped, increasing the network queue length may be advisable. Occasionally, a network analyzer will detect that the server has responded to an NFS request before the client retransmits the request. This indicates that the network queue on the NFS client system was full and the response was dropped. **ACTION 2:** *Increase maximum transmission unit* - The **nfsstat -in** command will display the maximum transmission unit (MTU) for each network interface active on the NFS client system. The MTU can be reset with the **ifconfig** command. This command is available on most UNIX NFS clients. For the z/OS NFS client, MTU can be reset in the TCP/IP profile. Increasing the MTU can improve overall performance and alleviate excessive IP fragmentation. However, with heterogeneous networks the MTU should not be set larger than the smallest acceptable MTU over the entire network route. For instance, when the network consists of a mix of Token Rings and Ethernets, the MTU is typically set at 1492 that is less than the 1500 MTU for an Ethernet and a 2000 MTU for Token Ring. Changing the network topology may allow you to increase the NFS client system MTU. **ACTION 3:** *Change socket buffer size* - Changing the buffer size is another way to tune the NFS client system. The AIX/6000 **no -a command** will display the current settings, in particular the **udp_sendspace** and **udp_recvspace**. The AIX/6000 **no -o udp_sendspace=32768** and **no -o udp_recvspace=32768** commands can be used to reset the **udp_sendspace** and **udp_recvspace**, respectively. For the z/OS NFS client, the **udpsendbfrsize** and **udprecvbfrsize** in the TCP/IP profile should be set respectively. The recommended value for Sun NFS V3 protocol is 65536. When processing with TCP protocol, it may be necessary to modify the **tcpsendbfrsize** and **tcprecvbfrsize**. **ACTION 4:** *Increase number of BIODs* - If data on an NFS server is typically accessed sequentially, or if there are many users on an NFS client system, it may be advantageous to increase the number of block I/O daemons (BIODs). This increases the NFS client processor utilization and should be weighed against other processor requirements on a multiple user system. If access is typically at random offsets within a file or file sizes are very small, you may even consider decreasing the number of BIODs. It should also be noted that some NFS client implementations honor BIODs differently to provide read or write bias. **ACTION 5:** *Increase timeout parameters* - If you have determined that NFS retransmissions are not caused by dropped NFS responses, you may consider increasing the timeout value. *This should be one of the last alternatives you select since this can have adverse effects.* If information is lost or dropped between the NFS client and the NFS server, increasing the timeout value can make performance worse. In this case the client would wait longer to decide to resend a request. In a heavily used NFS environment, increasing the timeout value also increases the likelihood that internal client or server buffers will become stagnant or stolen to service other network requests. You would probably only want to attempt this if you have strong evidence, probably from a network analyzer, that the NFS server is sending responses *after* the NFS client retransmits requests. The timeout values can be reset with the **mount** command.

NFS server system performance tuning

Regardless of whatever network and NFS client system tuning is done, some level of tuning or resource balancing should be addressed on the z/OS NFS server system. Typically, these systems have tuning and capacity planning procedures in place along with a high level of experience. Consequently, this section will be limited to a surface level discussion of the resource components to be considered when tuning the NFS server system.

z/OS constraints

Table 48 contains z/OS constraints information. These actions should be weighed against any increase in z/OS storage they may require.

Table 48. z/OS constraints symptom and action information

Symptom	Action
High CPU utilization with low aggregate throughput	**ACTION 1:** *Increase z/OS TCP/IP send and receive buffers* - There are other applications besides z/OS NFS which are dependent on the z/OS TCP/IP application. Some applications have different tuning requirements. The default values for **udpsendbfrsize** (or **tcpsendbfrsize**) and **udp udprcvbrfsize** (or **tcprcvbrfsize**) is 16 KB. It may be useful to increase these parameters up to 64 KB for better throughput with large files. **ACTION 2:** *Increase z/OS TCP/IP UDP queue* - Another area to consider is the z/OS TCP/IP UDP queue. As discussed in section "NFS client system performance tuning" on page 200, NFS responses are dropped when network adaptor queues are full. A similar situation occurs when the UDP queue of z/OS TCP/IP is full; in this case incoming NFS requests are dropped. The **noudpqueuelimit** keyword in the **assortedparms** section of the TCP/IP profile data set can be specified to enable the z/OS TCP/IP server to accept incoming UDP datagrams. Without specifying this keyword, the default queue length of 30 may not be sufficient. **ACTION 3:** *Modify NFS and TCP/IP dispatching priorities* - It is recommended that the z/OS NFS procedure have a relative dispatching priority less than that of TCP/IP. With MVS/ESA Version 5 Release 1.0 or later this is especially important since the System Resource Manager (SRM) algorithms have changed. Prior to this release of MVS mean time to wait algorithms resulted in the z/OS NFS receiving a lower dispatching priority than TCP/IP due to the order in which the address spaces were started. With this newer release of MVS mean time to wait dispatching priorities were adjusted based on increased I/O activity. Since TCP/IP has higher network I/O than z/OS NFS, the TCP/IP dispatching priority is lowered. Assigning fixed dispatching priorities, with TCP/IP dispatched at a higher relative value than the z/OS NFS, can alleviate this situation. **ACTION 4:** *Select transport protocols* - z/OS NFS supports TCP and UDP as transport protocols for the server, for both NFS Version 2 and Version 3 protocols. For the Version 4 protocol, NFS supports only TCP as a transport protocol. UDP is primarily used for its efficiency on high bandwidth, low latency networks, such as LANS. TCP is used for its efficiency on low bandwidth, high latency networks, such as WANS.
Constrained throughput with low CPU utilization	**ACTION 1:** - See "Subtasking" on page 217.

Table 48. z/OS constraints symptom and action information (continued)

Symptom	Action
High DASD utilization and NFS command response time	**ACTION 1:** *Evaluate placement of z/OS catalog data sets* - The performance of z/OS NFS will be impacted when serving files located on heavily utilized storage devices or devices behind congested storage controllers. Before deciding whether or not to move data or upgrade storage systems, make sure system catalogs are not on heavily utilized devices. Many of the NFS commands that are executed for conventional MVS data sets involve catalog access. In fact, listing directories, one of the most commonly executed commands on NFS client systems, accounts for a significant portion of the NFS **get** attribute and **lookup** commands. Such commands not only cause the z/OS catalog to be accessed but might also cause the entire file to be read depending on: • z/OS NFS processing options • Prior access via z/OS NFS • Whether or not files are DFSMS-managed **ACTION 2:** *DFSMS-managed data accessed from the network* - Files that are primarily accessed by way of z/OS NFS should be DFSMS-managed for improved performance. The z/OS NFS maintains file attributes for DFSMS-managed data sets when the z/OS NFS **nofastfilesize** processing option is specified. The **nofastfilesize** processing option provides exact file size determination rather than approximating file size as with the **fastfilesize** processing option. DFSMS management also provides improvements in terms of better storage utilization and specification of service levels. **ACTION 3:** *Evaluate placement of data accessed from the network* - While deciding whether or not to access DFSMS-managed files via the z/OS NFS, consider also the placement of such files. Files with critical performance requirements should be placed on the appropriate devices via the storage class parameters. If reallocation is necessary, you might also consider allocating sequential access method striped data sets for larger files or a z/OS UNIX file for smaller files. If you elect to maintain data set organization, you may choose to reblock the data set to a more optimal block size, such as half track blocking. You may even determine for some files that the best alternative is to store the files locally on the NFS client system.

Chapter 15. Evaluating z/OS NFS performance

This chapter assists both NFS client users and NFS server system administrators in evaluating the performance of z/OS NFS. Test methods are provided to assist in collecting performance data for evaluation. The appropriateness and usage of these test cases and methods depend on the following conditions:

- Workload requirements to be evaluated
- Level of expertise of the evaluator
- Monitoring tools available

Evaluating throughput

Throughput refers to the rate at which data is transferred between the NFS client and server systems. Establishing throughput baselines for the Client-Server environments of interest will assist in determining any benefits achieved from tuning. Baselines and measurement techniques are also useful in diagnosing performance problems.

As discussed in "How is performance characterized?" on page 193, a set of work activities or workloads should be identified for measurement. For throughput the type of workload selected should reflect the type of NFS read and write access expected by typical users or specific applications. Consider the following points:

- Average file size or range of file sizes
- Sequential or random access
- Data set types and organizations to be accessed
- Requirements for ASCII/EBCDIC translation

Single process throughput

The easiest place to begin evaluating NFS read and write throughput is on a single client and a single process basis. While measuring in a controlled and isolated laboratory situation is advantageous, it may not be practical. Measuring the actual environment of interest can provide more meaningful information as well as establishing a methodology for monitoring performance in the future. Under such uncontrollable situations, it is important to determine range and normal variation in measurements with particular attention to peak and non-peak periods of activity.

Once the NFS environment has been selected for measurement, some simple techniques can be used to initially evaluate NFS read and write throughput. One of the easiest methods is to use commands available on the NFS client system, such as **copy** or **cp**, to generate NFS reads and writes. Copy commands along with **date**, **time**, or **timex** commands can be used to determine the elapsed time for the copy. Throughput can be readily calculated from the elapsed time and the size of the file.

Some additional considerations when measuring NFS throughput are the effects of graphical user interfaces (GUI), any overhead associated with opening files, and I/O to and from local physical storage. Measuring with and without a GUI will show the effect on NFS throughput caused by the GUI. Using a network analyzer can assist with isolating the time spent opening files from the time spent executing NFS reads or writes. Techniques such as copying to /dev/null can eliminate local physical I/O for NFS reads. For more complex requirements such as reading and

© Copyright IBM Corp. 1991, 2006

writing at random offsets, you may use an existing application to provide such access or write a program to execute NFS procedure calls at random offsets within a file. Program generated NFS procedure calls can also be useful when available local physical storage is insufficient for the file sizes to be measured.

Multiple process throughput

For NFS client systems with multiple users or applications you may want to determine aggregate throughput on a multiple process basis. If throughput on a single process basis has achieved the network capacity, then multiple process throughput will be limited by the same network capacity as well as any NFS client system limitations. A simple method for evaluating multiple process throughput is to simply execute multiple single process measurements simultaneously.

Multiple client throughput

Evaluating multiple client throughput requires the ability to propagate the measurement workload over more than one NFS client system, preferably the same type of system. Such measurements are more complex in that they require controlled execution and collection of results from multiple systems. Depending on your network configuration, remotely executing your single process workload on multiple systems is a relatively simple method to use to initially evaluate multiple client aggregate throughput.

Evaluating NFS command response time

A user is probably more interested in the response time for typical user commands, such as listing a directory, making or removing a directory, and creating or removing a file. However, when such commands are executed for NFS mounted file systems, the NFS command response times become a factor in user command response time. The transparent nature of the NFS protocol allows a user to define a set of commands as a measurement workload with a minimum amount of effort.

For those desiring more detail on NFS command response time, there are public domain programs as well as industry standard benchmarks. Both of these options require more experience on the part of the user and may have limitations on their usage. Such programs are typically written for the UNIX environment and may not support all NFS client and server systems.

Evaluating CPU utilization

Evaluating CPU utilization on an MVS system can be quite complex, particularly for a production system. The activity on such systems is already generally monitored and z/OS NFS is simply one more application to be monitored. High level monitoring can be accomplished from the Display Active (DA) display of the System Display and Search Facility (SDSF). This display provides an overview of the total activity on the system. Of particular importance are the CPU usage and relative dispatching priorities of TCP/IP and z/OS NFS. Monitoring tools such as the Resource Measurement Facility (RMF™) Monitor I and Monitor II can provide more detail on the performance characteristics of the MVS system.

Collecting server usage data

You can use the SMF records that z/OS NFS produces to keep track of how files are accessed and how long each NFS user session lasts. The following SMF records can be produced:

Record type 42 subtype 7	This record, written when a file times out, provides file usage statistics.
Record type 42 subtype 8	This record, written when a client user logs out of the z/OS NFS, provides user session statistics.

Chapter 16. Tuning the z/OS NFS server

Previous chapters discussed performance tuning from a total system perspective. This chapter discusses performance tuning at the z/OS NFS server level. In particular, this chapter focuses on the z/OS NFS data set creation, processing, and site attributes.

Data set creation attributes

When creating new MVS data sets, the allocation parameters are the most readily accessible means for the end user to influence the performance of z/OS NFS server. They are also not well known outside of the z/OS user community. This is primarily due to the fact that MVS has a record oriented data structure and most NFS client systems have a byte oriented data structure. This section looks at how a variety of allocation parameters affect performance.

Block size and record length

Block size (BLKSIZE) specifies the maximum length, in bytes, of a physical block of storage in MVS. If BLKSIZE(0) is specified, the system will determine the optimal block size based on the maximum record length (LRECL) and the physical characteristics of the disk, or approximately half of a physical track. Half track blocking is optimal in that it reduces the amount of wasted storage on the disk.

The BLKSIZE is also important to the performance of the z/OS NFS, which performs physical I/O on a block basis. While half track blocking is generally optimal, for z/OS NFS server you should be aware of the relationship between the block size, the file record length, the network packet size and the NFS client **mount** write and read buffer sizes, *wsize* and *rsize*.

A network packet size will be negotiated between the NFS client and the NFS server when the NFS **mount** is processed. In general, the NFS server will determine the packet size for write operations, and the NFS client will determine the packet size for read operations. For z/OS NFS, this packet size is 8192 bytes, when processing with Sun NFS V2 protocol. For several NFS client implementations the default packet size is 8192. However, there are some implementations with a 4096 byte maximum packet size. The maximum packet size is 65536 bytes when processing with Sun NFS V3 protocol, and larger with the V4 protocol. However, z/OS NFS supports a maximum packet size of 32768 bytes for both the NFS V3 and V4 protocols.

The packet size can also be automatically reset to a smaller value when a constraint is detected. Such a reduction in packet size causes an increase in the number of NFS requests for the server to process the same amount of data. This increases the load on the server and may or may not have been the initial constraint causing the packet size reduction. This can have an adverse effect on z/OS NFS server performance, particularly when processing write requests.

Regardless of the packet size, the z/OS NFS server must reassemble the packets, possibly arriving out of sequence, into records. The z/OS NFS server buffers (see “Ordering out-of-sequence data” on page 215) and orders these packets until a

© Copyright IBM Corp. 1991, 2006

block of data can be written to disk. Write requests, especially those for large files, are processed more efficiently when the likelihood of packets spanning records and blocks is reduced.

When the z/OS NFS server is processing read requests, the physical I/O is again on a block basis. Consequently, read throughput is improved by increasing the block size up to half track blocking.

If files are allocated with a default block size (`BLKSIZE=0`) on RAMAC® 3 or Enterprise Storage Server (ESS™) 2105 DASD with 3390-3 format, the actual block sizes will probably vary between 24 and 27 KB, depending on the specified record length and record format. Table 49 shows the actual allocation block sizes for physical sequential data sets allocated on a RAMAC 3 or ESS 2105 DASD with a variety of record lengths and with both fixed and variable length records.

Table 49. Default block sizes for RAMAC 3 or ESS 2105 DASD

Record Format	Record Length (bytes)	Block Size (bytes)
FB	80	27,920
FB	1024	27,648
FB	4096	24,576
FB	8192	24,576
VB	84	27,998
VB	1028	27,648
VB	4100	27,998
VB	8196	27,998

Record format

Record format (RECFM) specifies the format and characteristics of the records in the data set. This section will be limited to a discussion of the performance considerations associated with fixed and variable length records.

There seems to be a natural affinity toward allocating MVS data sets with variable length records from the perspective of a byte oriented NFS client operating system. Such allocations enhance the feeling of transparent access and eliminate requirements to determine a fixed record length. However, the placement of incoming packets into records of varying lengths is more complex than placing packets into records with fixed lengths. Consequently, write performance for z/OS NFS is generally more efficient when MVS data sets are allocated with fixed length records.

Data set organization and data set type

Data set organization (DSORG) specifies the organization of an MVS data set. The z/OS NFS data set organization attribute can be physical sequential (ps), direct access (da), or VSAM. z/OS NFS server supports three types of VSAM files: key-sequenced (KSDS), entry-sequenced (ESDS), and relative record (RRDS). However, keyed access and relative-number access are not supported by the NFS protocol.

The DSORG attribute is ignored for directory-oriented NFS client commands. The data set type (DSNTYPE) specifies whether a PDSE or PDS is to be created when the z/OS NFS receives a **mkdir** command from the NFS client. A PDSE is created

when library is specified, and a PDS is created when PDS is specified. You cannot create another PDS (or PDSE) within a PDS (or PDSE).

Another MVS data set type is extended (EXT) which identifies an extended format data set. With z/OS NFS these DFSMS-managed data sets are allocated based on specification of a data class (DATACLAS). An appropriate data class must have been defined on the MVS system for an extended format data set to be allocated. One of the values specified when creating an MVS data class is the volume count which determines the number of stripes allocated for an extended format data set. If automatic class selection (ACS) routines have been written to allocate extended format data sets based on such criteria as naming conventions, it would not be necessary to specify a DATACLAS to allocate an extended format data set via the z/OS NFS server.

The hierarchical file system (HFS) and zSeries file system (zFS) are other data set types. They must be DFSMS-managed and be mounted within the z/OS UNIX subsystem. While an HFS or zFS file cannot be allocated using z/OS NFS, they can be accessed through the z/OS NFS server.

Data set organizations and data set types are generally selected based on user and application requirements. However, performance should be considered as one of these requirements. An important factor to consider when evaluating performance is the size of the files to be accessed. For smaller file sizes, we can differentiate between the overhead costs associated with creating or opening a file for output. As the file size increases, this impact becomes less of a factor in the transfer rates.

Processing attributes

Processing attributes are used to control how files are accessed from the NFS client system. Default values can be specified when the server is started. These defaulted values can be overridden at the NFS **mount** level or at the individual command level. The processing attributes, like the data set creation attributes, can affect the performance of the Network File System (NFS). This section will be limited to those processing attributes which have the most influence on performance.

Character translation

The processing attributes that specify whether or not data conversion occurs between ASCII and EBCDIC formats may have the most influence on the performance of z/OS NFS. The binary attribute, which specifies no data conversion, is the most efficient manner in which to access data by way of z/OS NFS. However, data conversion may be a requirement, particularly when data is shared with z/OS users or applications. Under such circumstances the text processing attribute would be specified instead of the binary processing attribute. The actual impact associated with data conversion will depend on the data to be converted.

File size determination

The attributes discussed previously have primarily affected data transfer rates. Whether or not to determine the exact file size in bytes has more of an effect on user command response time. This is because the file size in bytes is one of the file attributes obtained by an NFS **Get Attribute** procedure. The NFS **Get Attribute** procedure is executed for every user command accessing an NFS mounted file system.

z/OS NFS provides processing attributes to specify estimated file size determination, fastfilesize, or exact file size determination, nofastfilesize. The accuracy of the file size with fastfilesize processing is best for binary processing of data sets with fixed length records. For other situations, or when exact file size is a requirement, the **nofastfilesize** processing attribute should be specified.

When the **nofastfilesize** processing attribute is specified, there are performance factors to consider. The most important factor is that file size in bytes is maintained by z/OS NFS for DFSMS-managed data sets with many data set organizations.

Figure 30 compares the differences in response time between the z/OS NFS attribute support for DFSMS-managed data sets and the case when attributes are not cached for non-managed data sets. In this example, all data sets were members of a PDSE data set allocated with a variable block record format. The ls -l command was executed with the z/OS NFS **nofastfilesize** and **binary** processing attributes.

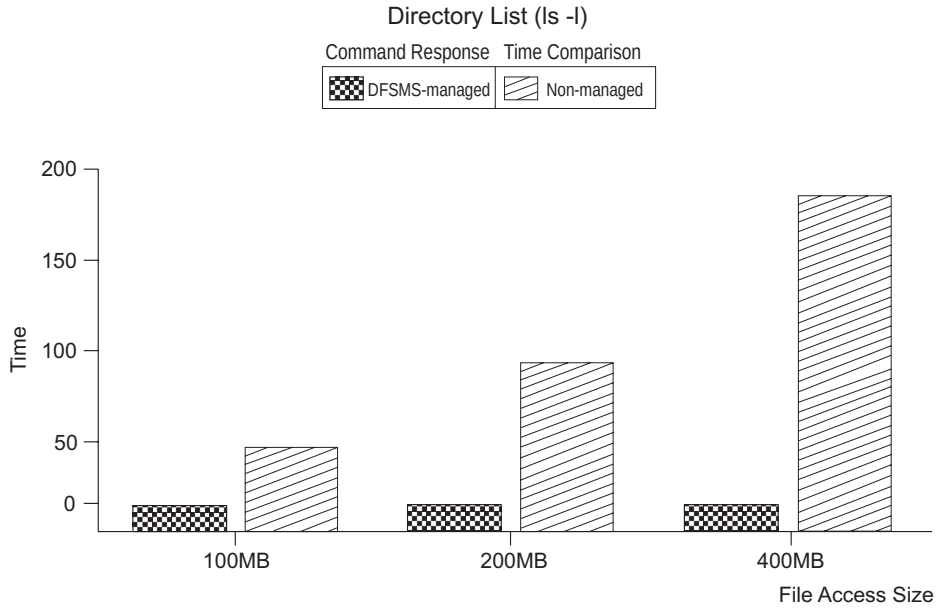

Figure 30. Directory list comparison between DFSMS-managed and non-managed

Data sets processed with the **nofastfilesize** attribute that are not DFSMS-managed will need to be read entirely to determine the exact file size in bytes when initially accessed by way of z/OS NFS. z/OS NFS will cache this information until the data set is modified (see section "Buffer usage and caching" on page 214) or the attribute time value has expired. The larger the file is, the greater the impact to command response time is. Figure 30 shows how the response time for a directory list command can be impacted when the file has to be read to determine the exact file size in bytes.

Maintaining the file size attribute by z/OS NFS also affects read throughput when accessing larger DFSMS-managed data sets with the **nofastfilesize** processing attribute. In the nofastfilesize case, reading the entire file to determine byte file size is not necessary. The impact on the end-to-end response time for the read is lessened by reducing the response time for the initial access to obtain file attributes.

Data set timeout specification

There are three timeout processing attributes that specify when z/OS NFS will release a data set following certain NFS operations. The **readtimeout** and **writetimeout** processing attributes specify how long z/OS NFS will wait after respective read and write operations before releasing a data set. The **attrtimeout** specifies how long a data set will remain allocated after an NFS **Get Attribute** or **Lookup** operation. When accessing data sets that are not DFSMS-managed with the **nofastfilesize** processing attribute, it may be advantageous to increase the attrtimeout value so that z/OS NFS caches data set attributes for a longer period of time. However, all three processing timeout attributes must be within the range of the **maxtimeout** and **mintimeout** site attributes.

Accessing migrated files

While migrating less frequently accessed data sets to tape improves MVS storage utilization, retrieving migrated data sets can impact the performance of z/OS NFS. z/OS NFS provides three processing attributes to specify how migrated data sets are to be handled by the server: **retrieve(wait)**, **retrieve(nowait)**, and **noretrieve**.

The **retrieve(wait)** attribute instructs the server to wait for a migrated data set to be recalled to a direct access storage device. The NFS client user or application process will not receive a response from z/OS NFS until the data set has been recalled. If the migrated file is relatively large, or migrated to tape, this can cause the NFS client to retransmit the request. The NFS client will retransmit the request a fixed number of times as specified by the **retrans** parameter of the **mount** command. If the recall of a migrated data set takes longer than the product of the timeo value on the **mount** command and one plus the **retrans** value for the **mount**, **timeo * (retrans +1)**, the initial command will most likely need to be executed again. Using the default values for the AIX/6000 **mount** command of seven-tenths of a second for **timeo** and three attempts for **retrans**, a recall would need to take less than 3 seconds. An alternative approach is to use the **retrieve(nowait)** processing attribute.

The **retrieve(nowait)** attribute instructs the server to recall a migrated data set and immediately return a ″device not available″ message to the client without waiting for the recall to complete. With this option, users attempting to access the file can continue to use their session for other activity rather than waiting an indeterminate time for the recall of a file. They can attempt to access the file again after allowing some period of time for the recall to have completed.

If it is critical to the user that a file be recalled before further work can continue, the **retrieve(wait)** processing attribute can be specified more selectively as part of the NFS client user command syntax along with the file name. While doing this provides the user more control over command execution, the user command may still timeout and have to be executed again.

Specifying the **retrieve(wait)** processing attribute as the system default may also impact the availability of z/OS NFS subtasks to process such requests. The last parameter of the **nfstasks** site attribute determines the number of z/OS NFS subtasks available to process such recall operations. See “Subtasking” on page 217 for more information about the **nfstasks** site attribute.

Another method of avoiding the unnecessary recall of migrated data sets is to specify the **noretrieve** processing attribute. With the **noretrieve** attribute, the server

does not recall a migrated data set and a ″device not available″ message is returned to the client. This processing attribute is particularly useful when listing a directory with unknown contents.

Asynchronous z/OS UNIX processing

z/OS NFS, with Sun NFS V2, protocol provides two processing attributes that only apply to z/OS UNIX processing. The **sync** attribute specifies that z/OS UNIX write requests should be committed to nonvolatile media (for instance, DASD) when received by the server. Some performance improvement can be obtained with the alternative **async** processing attribute. This type of processing causes z/OS UNIX write requests to be cached at a level within the z/OS UNIX architecture prior to the physical I/O. For more information about z/OS UNIX tuning, see “Ordering out-of-sequence data” on page 215 and “Subtasking” on page 217.

The **async** and **sync** attributes specified in the control file will be ignored when z/OS NFS server is communicating with Sun NFS V3 and V4 protocols. The processing attribute is determined in the client application implementation. Asynchronous write is also recommended for faster throughput.

Site attributes

The site attributes control z/OS NFS resources. Tuning of these parameters should be done with caution. They are only specified at the start of the z/OS NFS server and cannot be modified by client users. This section is limited to those attributes which more directly affect performance.

Buffer usage and caching

The **bufhigh** site attribute specifies the maximum size in bytes of allocated buffers before any buffer reclamation is initiated (see Figure 31). When the **bufhigh** limit has been reached, a percentage of the buffers will be reclaimed, and the amount of reclamation is determined by the **percentsteal** attribute. z/OS NFS uses this buffer area to cache file information, thereby satisfying requests more efficiently.

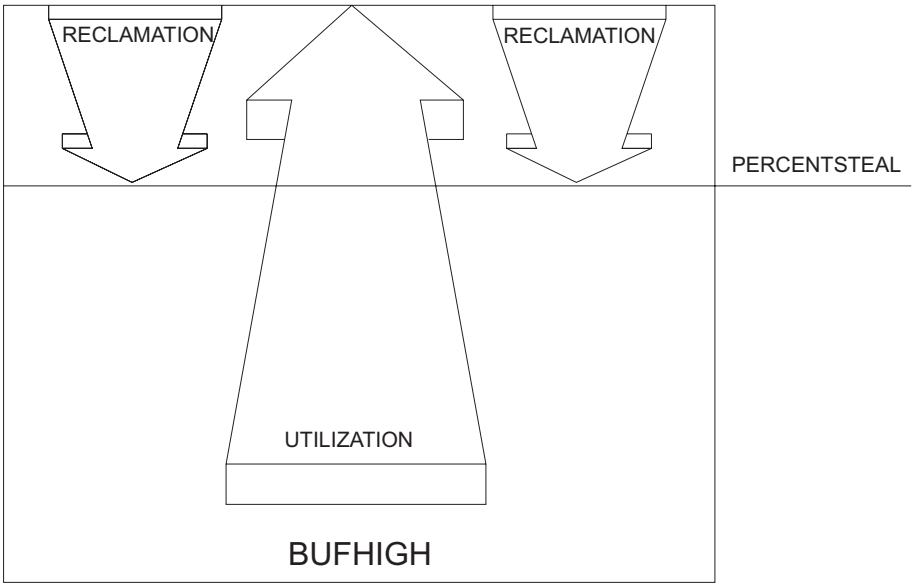

Figure 31. Bufhigh utilization with percentsteal

There are two additional site attributes, **readaheadmax** and **maxrdforszleft**, that also affect buffer usage. The **readaheadmax** attribute specifies the number of bytes to be read to fill the internal buffers so that additional read requests for that file may be satisfied directly from cache. The **maxrdforszleft** site attribute defines the number of physical block buffers to cache after determining a file's size (see section "File size determination" on page 211). This information is cached to satisfy any subsequent read requests for the same file.

Keep in mind also that some information will be cached on a file handle basis, in other words, for every file accessed within the timeout periods. Tuning of the **bufhigh**, **readaheadmax**, and **percentsteal** values should be determined based on the following conditions:

- Number of files to be accessed
- Available region
- Amount of data to cache per file

For example, suppose that, on average, 1,000 physical sequential (DSORG=PS) files allocated on 3390 Model 3 DASD are accessed within the same timeout period via the z/OS NFS server. Further, assume that we would like two blocks of a file cached internally to satisfy read requests. If we assume an average block size of 25 KB (see "NFS client system performance tuning" on page 200 for more detail), we would need at least 50 MB of storage for internal buffers. With a 20 percent value for **percentsteal**, a 64 MB value for **bufhigh** would probably be more reasonable. For this example, a reasonable value for the **readaheadmax** attribute is the file block size, or 25 KB. The **readaheadmax** value should not be less than twice the maximum record length of files accessed by way of the z/OS NFS server.

Ordering out-of-sequence data

While the previous section primarily addressed satisfying read requests from cache, this section discusses the caching used by z/OS NFS to satisfy write requests. z/OS NFS provides two site attributes for this purpose: **cachewindow** and **logicalcache**. As with all site attributes, **cachewindow** and **logicalcache** can only be modified at server startup.

The **cachewindow** attribute specifies the window size, in terms of number of packets, used in logical I/O to buffer client block writes received out of sequence. Figure 32 on page 216 shows how out of order packets are buffered in a cache window until a block of data can be physically written to storage. Since out of sequence packets are an inevitability, a certain amount of buffering is necessary.

In a congested network environment, the nature of the underlying protocols should be taken into consideration when tuning the **cachewindow** attribute. Under such circumstances, NFS, UDP, and IP requests are more likely to timeout. When this situation is detected within the underlying protocols, the packet size can be reduced automatically. The packet size is generally halved until the server and client systems are balanced. The minimum packet size that can be negotiated is 512 bytes. Accessing the z/OS NFS server in such an environment can result in the default 8 KB read and write buffers being reduced to 512 bytes, or 16 packets instead of one packet. *Under these circumstances the recommended value for* **cachewindow** *is 16 times the number of BIODs.*

Another factor to consider in tuning the **cachewindow** attribute relates to the relationship between packet size, data set record length, and block size. If the packet size is typically greater than or equal to the data set block size, there is less of a need to buffer the packets. On the other hand, if the packet size is smaller

than the block size and/or record length, a larger **cachewindow** value will be required to buffer the packets until a block I/O is executed.

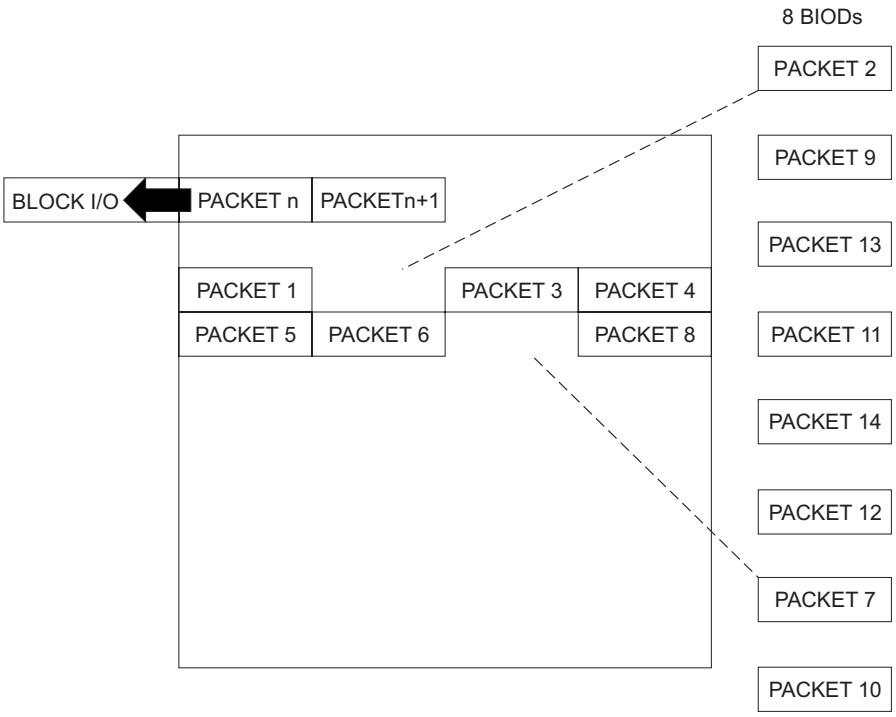

Figure 32. Relationship between cachewindow and BIODs

The **logicalcache** site attribute specifies the size (in bytes) of allocated buffers for all the cache windows combined (see Figure 33). The number of cache windows in use will depend on the number of file handles accessed within a given timeout period.

LOGICALCACHE

CACHEWINDOW			CACHEWINDOW		
PACKET	PACKET	PACKET	PACKET	PACKET	PACKET
PACKET		PACKET	PACKET	PACKET	
	PACKET	PACKET	PACKET		PACKET
PACKET				PACKET	
CACHEWINDOW			CACHEWINDOW		
PACKET	PACKET	PACKET	PACKET	PACKET	PACKET
PACKET	PACKET		PACKET		PACKET
	PACKET	PACKET	PACKET	PACKET	PACKET
	PACKET		PACKET		

Figure 33. Logicalcache utilization for cachewindows

Tuning of the **logicalcache** and **cachewindow** values should be determined based on the following conditions:

- Number of files accessed for write
- Available region
- Typical packet size
- Data set block size
- Number of client Biods

The **cachewindow** and **logicalcache** attributes do not apply to z/OS UNIX processing.

For example, suppose you have AIX NFS client workstations on a lightly loaded network, with the default value of six Biods, writing to physical sequential MVS data sets allocated with 8 KB fixed length records. Further, assume that these files were allocated on 3390 Model 3 DASD with a 24 KB block size and that the network packet size is 8 KB. Then, a value of 12, twice the number of NFS client Biods, for **cachewindow** means that an out of sequence packet could be internally buffered within a 96 KB or four block range. If a packet is received outside this range, z/OS NFS must read data from storage in order to write the packet at the correct offset into the data set. z/OS NFS detects and ignores duplicate requests.

A value of 12 for **cachewindow** and a network packet size of 8 KB in our example results in a 96 KB buffer requirement per client actively writing to MVS data sets via z/OS NFS. A 5 MB value for **logicalcache** is sufficient to satisfy sustained write requests from 50 NFS client workstations at the same time.

Storage considerations

To avoid storage shortages, consider the following conditions:

- If a region size is specified, the **bufhigh + logicalcache** value should not be more than 90 percent of the region size. (The more **bufhigh**; the better performance.) With **bufhigh(1M)**, **logicalcache(1M)**, and the existing NFS core size, the minimum region size is 25M.
- If a region size is equal to 0, **bufhigh + logicalcache** value is limited by the supported virtual memory size on the installation. Region size can be set up as large as the size of the local page data sets.

The System Programmer should also consider if the other jobs on the installation use virtual memory.

Subtasking

The **nfstasks(n,m,o)** site attribute specifies the number of server processes to initiate on startup. If **nfstasks(n,m)** is specified, then the valid value range for both n and m is 1 to 100. The sum of n and m must be less than or equal to 101. If **nfstasks(n,m,o)** is specified, then the valid value range for 'n' and 'm' is from 1 to 100. The valid value range for 'o' is from 1 to 101. The sum of 'n' plus 'o' must be less than or equal to 101.

The absolute and relative value of *n* and *m* should be tuned for the expected system usage. If conventional MVS data sets will be accessed, primarily, then *n* should be relatively high. If z/OS UNIX files will be accessed, primarily, then *m* should be relatively high. The absolute value of these will influence the amount of system resources consumed (higher values will make more system resources available to process NFS requests.

The parameter *n* is the number of processes started to handle asynchronous operations (conventional MVS data set access) and short duration synchronous operations (SAF calls). The parameter *m* is the number of processes which handle z/OS UNIX requests. Increase this value if your server requires many parallel z/OS UNIX accesses. The parameter *o* is the number of processes which handle longer duration synchronous operations (concurrent NFS server recall requests.) Increase this value if your server supports lots of active recall operations. When **nfstasks(n,m)** is specified, m represents both the processes for z/OS UNIX as well as the long duration synchronous operations. It is better to use the **nfstasks(n,m,o)** format to have better granularity of control as well as have the capability to specify a higher number of z/OS UNIX processes to handle more z/OS UNIX operations.

Chapter 17. Tuning the z/OS NFS client

This chapter discusses performance tuning at the z/OS NFS client level. Performance of the z/OS NFS client is affected by the following installation parameters that can be tuned (see Table 50). The default value of each parameter is specified in parentheses.

Table 50. Client installation parameters for tuning
attrcaching (Y)
biod (6)
bufhigh (32)
datacaching (Y)
delaywrite (16)
dynamicsizeadj (Y)
readahead (8)

Most of the installation parameters that are specified in member BPXPRM*xx* in SYS1.PARMLIB can be overridden for a mount point by using the **mount** command. The exception to this is **biod** and **bufhigh**, which require restarting z/OS UNIX.

Table 51 shows mount parameters that affect performance and that can be tuned. The default values are specified in parentheses.

Table 51. Mount parameters for performance and tuning
attrcaching (Y)
datacaching (Y)
delaywrite (16)
readahead (8)
rsize (8192 for V2) (No default for V3 - system determined)
wsize (8192 for V2) (No default for V3 - system determined)
vers (3)

Caching

Caching of file attributes (**attrcaching=y**) and data (**datacaching=y**) are recommended for performance. By caching the file data, all subsequent references to the cached data is done locally, avoiding the network overhead. Further authorization checking for subsequent access to the cached data or for other client users is done on the z/OS NFS client system. With attribute caching, z/OS NFS client will perform consistency checks to maintain valid cached data.

Dynamicsizeadj

The **dynamicsizeadj** parameter specifies the capability to have the packet size adjusted automatically when the remote procedure call (RPC) timeout has occurred.

For example, the **mount** command specifies, rsize=32K. If the RPC READ with a 32 K bytes request times out:
- The RPC request times out.
- The client cuts the size in half and sends the RPC READ with a 16 K request.

© Copyright IBM Corp. 1991, 2006

- The RPC request times out again.
- The client cuts the size in half and sends the RPC READ with a 8 K request.

Note: This process continues until the minimum size is reached.

Bufhigh

The **bufhigh** parameter specifies the maximum size in megabytes of allocated buffers for caching. z/OS NFS client uses this buffer area to cache file data, thereby satisfying requests efficiently.

When a client system is suffering high CPU utilization and low throughput, increase the **bufhigh** storage size. Increasing the **bufhigh** storage size by twice the default value (within available region) may improve system performance. Also, if the client system is processing mostly large files, it is good to increase the **bufhigh** storage size. Any changes to the **bufhigh** parameter require restarting z/OS UNIX.

Biod

The z/OS NFS client uses block I/O daemons (BIOD) to perform asynchronous read or write when **datacaching** is enabled. The **biod** parameter specifies the number of asynchronous block I/O daemons to use. If there are a lot of users on the NFS client system, it may be advantageous to increase the number of BIODs. The maximum number of BIODs should never exceed two times the number of client system processors. This will avoid excessive TCB switching. Tuning of this parameter should be done with caution. Any change to the **biod** parameter requires restarting z/OS UNIX.

Readahead

With caching enabled, read requests may be satisfied from cache. The **readahead** parameter specifies the maximum number of disk blocks to read ahead. Thus, additional read requests for the file may be satisfied directly from cache. When application is anticipating mostly sequential reads, increase the readahead to the maximum. If application is mostly processing random reads, use the default or a minimal readahead value.

Delaywrite

With caching enabled, write requests may be cached. The **delaywrite** parameter specifies the maximum number of disk blocks to cache. The z/OS NFS client will issue the WRITE RPC whenever the system load is low or if the cache data are in sequence. **delaywrite** increases WRITE performance by overlapping the disk I/O. This performance increase is most apparent with many small writes or with many random writes. When there is a high volume of write requests, set the **delaywrite** parameter to the maximum value(32). If there is a low volume of write requests, set the **delaywrite** to the default (16).

Vers

Sun NFS version 3 protocol support is available on OS/390 V2R6 NFS. It has performance enhancements over Sun NFS version 2 protocol. The parameter does not have to be specified unless you would like to override the default. The z/OS NFS client will communicate at the highest protocol level supported by the server, by default. Verify the protocol level for tuning of read or write buffer size (**wsize**,

rsize). Use nfsstat –m /mountpoint command and look for **vers(n)** in the report or use rpcinfo –p servername command and look for 100003 3 udp 2049 nfs in the report.

Wsize and rsize

wsize and **rsize** specify the buffer size to use for read and write request. For NFS Version 2 protocol, **rsize** and **wsize** are a multiple of 512 bytes, up to a maximum of 8192 bytes.

For the NFS server with Version 3 protocol, **rsize** and **wsize** are negotiated between the z/OS NFS client and the NFS server. The maximum buffer size that is supported in z/OS NFS client system is 32 KB. The negotiated read and write buffer size will be the smaller of the buffer size supported by the server and 32 KB (min(server_info, 32KB)).

The z/OS NFS client system will reduce the read or write buffer size when there is constraint on the storage resource.

When using NFS Version 3 protocol the system normally determines the optimal buffer values. However, if the number of retransmissions or the number of time-out from **nfsstat** report is high, reduce the write buffer size (**wsize**) to 8 KB to avoid retransmission of large 32 KB data.

Part 5. Diagnosis and Messages

Chapter 18. Diagnosis and reporting of problems 225
Correcting input errors 225
Using keywords to identify a problem 226
Component identification keyword 226
Release level keyword 226
Type-of-failure keyword 227
ABENDxxx 227
MSGGFShnnnt 227
WAIT 227
LOOP 228
INCORROUT 228
DOC SCnnnnnnnn 228
PERFM 229
Service level keyword 229
| Using z/OS component tracing 229
| Component trace benefits 230
| Using NFS component trace PARMLIB member
| CTINFS*nn* 230
| CTINFS00 member of SYS1.PARMLIB . . . 231
| Capturing NFS component trace information in
| an SVC dump 232
| Using a z/OS component trace external writer 233
| Step 1. Telling MVS to start a CT external
| writer 233
| Step 2. Telling NFS to connect to a CT
| external writer and start sending records . . 234
| Using IPCS to view records from an external
| writer 234
Setting up a dump data set for abnormal ends . . 235
Searching the IBM database for APARs and PTFs 235
Contacting the IBM Support Center 235
Diagnostic aids 236
First failure data capture 236
Symptom data 236
SVC dump 236
Dump suppression 237
Data capture suppression 237
Invocation 237
Errors and messages 237
Debug trace data capture 238
Environmental checklist 238

Chapter 19. Network File System messages . . 241
Server messages 241
Client messages 281
Messages from the client platform (AIX) 295

Chapter 20. Return codes 299

| **Chapter 21. Reason codes** 309

© Copyright IBM Corp. 1991, 2006

Chapter 18. Diagnosis and reporting of problems

This chapter is intended to help you diagnose NFS problems and use keywords to describe the NFS program failures.

Before you begin diagnostic procedures, verify that the suspected problem is not a result of failure caused by:

- z/OS Communication Server's TCP/IP (Transmission Control Protocol/Internet Protocol)
- Broken clients (for example, sending incorrect file handles or ignoring the server's error return code)
- The network (for example, packets not arriving at the server, clients not receiving replies, or a gateway going down)
- The result of user input error.

For failures caused by z/OS Communication Server's TCP/IP, see *z/OS Communications Server: IP Diagnosis Guide* to diagnose the problem.

Correcting input errors

For a user input error, use this procedure to assist in correcting the error:

1. Look up the command and its attributes (or parameters), the attributes in the attributes data set, and the entries in the exports data set in this book or the appropriate client machine command documentation.
2. Examine the attributes (or parameters) specified by each command, specified in the attributes data set, and the DIRECTORY entries specified in the exports data set to verify that they are specified correctly. Check the z/OS NFS server log data set, the z/OS NFS client log data set, or both log data sets to find the error message if applicable.
3. If you find a user input error, reenter the command or restart the server after you correct the error.
4. If all attributes (or parameters) and directory statements appear to be specified correctly, you should use this chapter to build a set of keywords that describe the error, and then contact IBM for help.

After determining that the suspected failure is neither a z/OS TCP/IP program error nor a user error, you need to develop a set of keywords that describe the program failure. A keyword is an agreed-upon word or abbreviation used to describe a single aspect of a program failure.

After you have selected a set of keywords, use the set to search IBMLink™/Service to determine whether an authorized program analysis report (APAR) has already been recorded for the failure. The IBMLink/Service online database contains information about the resolution of APARs. If a program failure is identified in these databases with the same set of keywords, the search yields a description of the problem and usually a fix. If the failure is not known, use the keywords to describe the failure when you contact IBM for help, which is the final step before an APAR is generated. For more information about these services see *z/OS DFSMSdfp Diagnosis*.

© Copyright IBM Corp. 1991, 2006

Using keywords to identify a problem

In general, when you contact IBM, you are asked to identify the problem with a full set of keywords. A full set of keywords for the NFS identifies these specific keywords:

- Component identification keyword
- Release level keyword
- Type-of-failure keyword
- Service level keyword

This example displays a full set of keywords:

```
5695DF121 R1VS MSGGFSA865I UW12345
```

Notes:

1. 5695DF121 - Signifies the component identification keyword.
2. R1VS - R1VS is the release level keyword for the z/OS NFS server, FMID HDZ11VS. R1VC is the release level keyword for the z/OS NFS client, FMID HDZ11VC.
3. MSGGFSA865I - MSGGFSA is the type-of-failure keyword prefix for the z/OS NFS server. MSGGFSC is the type-of-failure keyword prefix for the z/OS NFS client.
4. UW12345 - Service level keyword.

These sections explain individual keywords and their relation to the full set of keywords used to describe a NFS program failure.

A search in IBMLink/Service using the first keyword, the NFS component identifier, alone would detect all reported problems for the entire program product. Each keyword added to the search argument makes the search more specific, thereby reducing the number of problem descriptions to be considered. Sometimes you can find a correction for a problem without using the full set of keywords. Each keyword after the first is optional.

Component identification keyword

The component identification number is the first keyword in a set. You should use this keyword whenever the Network File System is suspected of being the failing component. The component identification keyword should be used with at least one other keyword to search IBMLink/Service. Used alone, it produces a full listing of APARs against the Network File System.

1. The component identification number for the Network File System is 5695DF121.
2. Go to "Release level keyword."

Release level keyword

Using this keyword to identify the release level of the Network File System is optional in the IBMLink/Service search argument. However, it is required when you submit an APAR.

1. Issue the MODIFY mvsnfs VERSION command, as shown in "Version operand" on page 172.
2. In output message GFSA944I, find the release level, as shown by *R1xx* in the following example:

   ```
   GFSA944I z/OS Network File System Server release R1xx ....
   ```

3. See “Type-of-failure keyword.”

Type-of-failure keyword

To identify the type of failure that occurred, select the one keyword from Table 52 which best describes the problem. Then follow the specific instructions for that keyword. If you are not certain which of two keywords to use for the type of failure, use the one that appears first in the list.

Table 52. Summary of type-of-failure keywords

Keyword	Type of Failure	Page
ABEND*xxx*	Network File System ends abnormally because of a system-detected error	227
MSGGFS*hnnnt*	Error indicated by a system message	227
WAIT	The program does not seem to be doing anything	227
LOOP	The program is doing something repetitively	228
INCORROUT	Output from the program is incorrect or missing	228
DOC SCnnnnnnnn	Documentation is incorrect or incomplete	228
PERFM	Performance of the program is degraded	229

ABENDxxx

Use this procedure when Network File System ends abnormally.

In this step, you build a type-of-failure keyword:

1. Replace the *xxx* part of the **abend*xxx*** keyword with the abend code from either the SYSLOG message or the abend dump. For example, if the abend code is 0C4, the ABENDxxx keyword is ABEND0C4.
2. To continue, see “Service level keyword” on page 229 to identify the service level of Network File System.

MSGGFShnnnt

Use this procedure for any of these conditions:

- Network File System message indicates an unexpected error detected.
- Message is not issued under conditions that should cause it to be issued.
- Message is issued under conditions that should not cause it to be issued.

In this step, you build a type-of-failure keyword.

1. For the *hnnnt* portion of the MSGGFShnnnt keyword replace the *h* with an **A** if the message prefix was GFSA or with a **C** if the message prefix was GFSC. Replace the *nnn* with the message number and replace the *t* with the type code. For example, if the message is GFSA865I, the MSGGFShnnnt keyword is MSGGFSA865I.
2. To continue, see “Service level keyword” on page 229 to identify the service level of the Network File System.

WAIT

Use this procedure when Network File System suspends activity while waiting for some condition to be satisfied without issuing a message to tell why it is waiting.

In this step, you build a type-of-failure keyword:

1. If a Network File System task was in a WAIT, use the ***wait*** keyword.

```
5695DF121 Rccc WAIT
```

2. To continue, see “Service level keyword” on page 229 to identify the service level of the Network File System.

LOOP

Use this procedure when some part of the program repeats endlessly.

In this step, you build a type-of-failure keyword.

1. If a Network File System task was in a LOOP, use the ***loop*** keyword.
2. If a message repeats endlessly, use the MSGGFS*hnnnt* keyword at the same time with the LOOP keyword. For the *hnnnt* portion of the MSGGFShnnnt keyword replace the *h* with an A if the message prefix was GFSA or with a C if the message prefix was GFSC. Replace the *nnn* with the message number and replace the *t* with the type code.

   ```
   5695DF121 Rccc LOOP MSGGFShnnnt
   ```

3. To continue, see “Service level keyword” on page 229 to identify the service level of the Network File System.

INCORROUT

Do not use ***incorrout*** if another keyword applies. Use this procedure only for these conditions:

- Output was expected but not received (missing).
- Output is different from expected (incorrect).

In this step, you build a type-of-failure keyword.

1. Use the incorrout keyword.
2. If the output is in the form of an incorrect message, also use the type-of-failure keyword MSGGFS*hnnnt*. For the *hnnnt* portion of the MSGGFShnnnt keyword replace the *h* with an A if the message prefix was GFSA or with a C if the message prefix was GFSC. Replace the *nnn* with the message number and replace the *t* with the type code.

   ```
   5695DF121 Rccc INCORROUT MSGGFShnnnt
   ```

3. To continue, see “Service level keyword” on page 229 to identify the service level of the Network File System.

DOC SCnnnnnnnn

Use this procedure when a program problem appears to be caused by incorrect or missing information in the Network File System documentation. If the documentation error is minor (for example, incorrect punctuation or a misspelled word), you do not need to build a keyword string to describe it. Instead, submit a Reader’s Comment Form from the back of the book.

If the error is serious, and of general concern to other users, continue with this procedure.

In this step, you build a type-of-failure keyword, which includes a document order number. No other keywords are needed:

1. Place the order number of the document after the DOC keyword, omitting the hyphens. If the suffix is one digit, precede it with a zero. For example, the keyword for a document with order number SC26-7029-01 would be DOC SC26702901.
2. Locate the page in the document at which the error occurs, and prepare a description of the problem. If you submit an APAR, you must include this information in the error description.

3. To continue, see "Service level keyword" to identify the service level of the Network File System.

PERFM

Performance problems can be related to system tuning. Use this procedure when the performance problem cannot be corrected by system tuning and performance is below stated expectations.

In this step, you build a type-of-failure keyword.

1. Use the ***perfm*** keyword.

   ```
   5695DF121 Rccc PERFM
   ```

2. To continue, see "Service level keyword" to identify the service level of the Network File System.

Service level keyword

Use this keyword to identify the service level of the module that failed. The service level is defined as the most current fix applied to the module.

1. Issue the MODIFY mvsnfs VERSION=module command, as shown in "Version operand" on page 172.
2. In output message GFSA945I, find the service level of the module, as shown by *service_level* in the following example:

   ```
   GFSA945I modulename service_level compile_date/time
   ```

3. Specify the service level keyword.

   For example, if the service level of the failure related module is UW12345, you would specify UW12345 as the service level keyword.

   ```
   5695DF121 Rccc MSGGFSA865I UW12345
   ```

You now have all the necessary information for an effective search of known problems in IBMLink/Service. Refer to "Contacting the IBM Support Center" on page 235 or "Searching the IBM database for APARs and PTFs" on page 235.

Using z/OS component tracing

To improve the recording and tracing of NFS diagnostic information, you can store diagnostic information for the z/OS NFS server in a z/OS component trace buffer. After activating component tracing and recreating the problem, the diagnostic information in the z/OS component trace buffers can then be captured in a z/OS SVC dump, written to a DASD or tape data set, and viewed using the IPCS CTRACE command. Use of component trace buffers offers other advantages over the log data sets (NFSLOG1 and NFSLOG2) including performance improvements and virtual storage constraint relief.

To activate component tracing of NFS information, an operator can issue a TRACE CT command specifying the procedure name of the z/OS NFS server with individual trace options or a CTINFS*nn* member of SYS1.PARMLIB that contains the trace options. The operator can also activate component tracing at NFS server startup, by specifying CTRACE=*nn* on a START mvsnfs command, where *nn* is the suffix of a CTINFS*nn* PARMLIB member. For details on these commands and their NFS trace options, see Chapter 11, "Network File System operation," on page 157. For details on the CTINFSnn member of SYS1.PARMLIB, see "Using NFS component trace PARMLIB member CTINFS*nn*" on page 230.

The following sections describe how to use NFS component tracing in diagnosis procedures. Further information on using the IPCS CTRACE command can be found in *z/OS MVS IPCS Commands*. Additional information regarding the MVS DISPLAY, TRACE, and DUMP commands may be found in *z/OS MVS System Commands*.

When you use component tracing, NFS diagnosis information continues to be recorded in the NFS log data sets, at the level specified on a MODIFY *mvsnfs*,LOG=*xxx* operator command. IBM recommends that you adjust the level of logging to avoid duplicating diagnostic records and to maximize the performance benefits of using component trace buffers.

If additional NFS diagnosis information is required using the log data sets for either the z/OS NFS server or client, see Appendix L, "Capturing diagnostic information using z/OS NFS log data sets and from other components," on page 379 for further details.

Component trace benefits

Component trace support offers the following benefits over the NFS log data sets.

- For virtual storage constraint relief, trace buffers are stored to a z/OS data space. Each z/OS NFS server in a multi-server implementation has its own data space initialized at server startup.
- Performance is improved because in-core tracing can now be activated without the overhead of I/O.
- The optional use of an external writer data set may extend the amount of trace data collected when accompanied with a z/OS dump.
- For performance, remaining trace data formatting routines have been removed from the clients' request mainline path and are now performed asynchronously by the error log daemon. The I/O overhead of writing to a data set is performed under the started external writer address space. Removing the I/O path provides a more responsive NFS error log daemon.
- Use of the IPCS MERGE command provides the ability to merge an NFS server component trace with those of other NFS servers and related z/OS products such as TCPIP and UNIX System Services to obtain a single image of system activity.
- Performance is improved over the previous NFS server tracing methodology with the ability to enable only those trace events of interest. Previously, the amount of tracing was cumulative.

Using NFS component trace PARMLIB member CTINFS*nn*

The CTINFS00 member of SYS1.PARMLIB (Figure 34 on page 232) provides default component trace values for the z/OS NFS server. This member can be invoked on the start command when the NFS server is started, or on a TRACE CT command while the server is running. You can create and use additional copies of the member with different trace option values, placing them in SYS1.PARMLIB using the naming convention CTINFS*nn*.

To start a component trace when the NFS server is active, the operator issues a TRACE CT command on the master console or a console with master authority. The syntax and options for a component trace of the z/OS NFS server are shown in "Starting component tracing for the z/OS NFS server" on page 161. The operator can specify the trace options individually on the command, or can specify the name of a CTINFS*nn* PARMLIB member containing the desired trace options.

Using a PARMLIB member on the TRACE CT command can help minimize operator intervention and avoid syntax or keystroke errors.

If you need to diagnose problems before the z/OS NFS server is fully active, component trace can be invoked on the START *mvsnfs* command using the CTRACE=*nn* parameter, where *nn* is the suffix of a CTINFS*nn* PARMLIB member. The component trace options specified in that member then control the component trace. These trace options can be overridden by a subsequent TRACE CT command if necessary, for example if a problem arises requiring a different set of trace options.

CTINFS00 member of SYS1.PARMLIB

Figure 34 on page 232 shows a copy of the CTINFS00 member of SYS1.PARMLIB that NFS provides for tracing the z/OS NFS server. This file is used for setting the initial startup trace settings. In the file:

- ON/OFF specifies whether trace is turned on or off.
- OPTIONS specifies the trace options that are to be applied.

The settings in CTINFS00 turn on these trace options: Error, Warning, Informational, and First Failure Data Capture.

```
/* ============================================================== */
/*                                                                */
/* $MAC(CTINFS00)  COMP(5694DF121)                              @P01C*/
/*                                                                */
/* Z/OS Network File System Server.                               */
/* Sample CTRACE options.                                         */
/*                                                                */
/* COPYRIGHT:                                                     */
/*PROPRIETARY V3 STATEMENT                                        */
/*Licensed Materials - Property of IBM                            */
/*"Restricted Materials of IBM"                                   */
/*5694-A01                                                     @P01C*/
/*(C) Copyright IBM Corp. 2004                                    */
/*END PROPRIETARY V3 STATEMENT                                    */
/*                                                                */
/*                                                                */
/* ---------------------------------------------------------------- */
/*                                                                */
/* CHANGE ACTIVITY:                                               */
/*                                                                */
/* $L77=NFS,HDZ11VS,040507,IBSMVB: Convert To CTrace           @L77A*/
/* $L77=NFS,HDZ11VS,041210,IBSKYL: Add FFDC                   @L777A*/
/* $P01=KAJ0199,HDZ11VS,041216,SJPLMB: Restore PID to 5694-A01 @P01A*/
/*                                                                */
/* ---------------------------------------------------------------- */
/* DEFAULT CTINFS00 MEMBER                                        */
/* ============================================================== */
/*                                                                */
TRACEOPTS
/* ---------------------------------------------------------------- */
/*   Optionally start external writer in this file (use both      */
/*   WTRSTART and WTR with same wtr_procedure)                    */
/* ---------------------------------------------------------------- */
/*       WTRSTART(wtr_procedure)                                  */
/*       WTRSTART(CTWTR)                                          */
/* ---------------------------------------------------------------- */
/*   ON OR OFF: PICK 1                                            */
/* ---------------------------------------------------------------- */
         ON
/*       OFF                                                      */
/* ---------------------------------------------------------------- */
/*       WTR(CTWTR)                                               */
/* ---------------------------------------------------------------- */
/*   OPTIONS: NAMES OF FUNCTIONS TO BE TRACED.                    */
/* ---------------------------------------------------------------- */
         OPTIONS(
                 'INFO'
                ,'WARNING'
                ,'ERROR'
                ,'FFDC'
                 )
/*                                                                */
```

Figure 34. z/OS NFS Server component trace PARMLIB member CTINFS00

Capturing NFS component trace information in an SVC dump

You can use component tracing to gather information for diagnosis when recreating prior problems related to the z/OS NFS server. To do so, follow these steps:

1. Have an operator issue a TRACE CT command, as described in "Starting component tracing for the z/OS NFS server" on page 161. Specify individual trace options on the command, or OPTIONS=All to include all trace types.
2. Recreate the reported problem. After the problem is recreated, disable tracing with the following command at the operator console:

   ```
   TRACE CT,OFF,COMP=mvsnfs
   ```

   Where *mvsnfs* is the name of the procedure in your system that PROCLIB used to start the NFS server.
3. Create an MVS dump

   To create a dump of the an NFS server address space and associated data space, enter the following command at the operator console:

   ```
   DUMP COMM=(any dump description title you choose)
   ```

   In response to operator message *nn IEE094D SPECIFY OPERAND(S) FOR DUMP COMMAND, reply with:

   ```
   R nn,JOBNAME=MVSNFS,DSPNAME='MVSNFS'.MVSNFS,CONT
   R nn,SDATA=(LPA,CSA,ALLNUC,GRSQ,LSQA,SWA,PSA,SQA,TRT,RGN,SUM)
   ```

   To include NFS component trace records in an SVC dump, specify the associated data space name *'MVSNFS'.MVSNFS* on the MVS DUMP command. The value of *MVSNFS* represents the server's procedure name.

   You can now use the IPCS command CTRACE COMP(MVSNFS) FULL to look at the trace records, where the default data set has been set to a captured dump data set. NFS provides both IPCS exits and format tables to facilitate this activity.

Using a z/OS component trace external writer

z/OS Component Trace (CT) supports an external writer which can be used to write trace records to a DASD or tape data set in real time. That is, as NFS generates trace records those records can be placed on DASD or tape.

Activating a CT External Writer for a program is a two-step process. These steps can be done in either order, but the external writer will not start writing out trace records until they have both been done.

Step 1. Telling MVS to start a CT external writer

Before a CT external writer can start there must be a procedure in SYS1.PROCLIB that tells how to invoke the writer. An example of a CT external writer procedure that can be placed in SYS1.PROCLIB(CTWTR) is:

```
//CTWTR    PROC
//IEFPROC  EXEC PGM=ITTTRCWR
//TRCOUT01 DD   DSN=IBMUSER.CT1,DISP=OLD
```

If you are tracing in a multi-server or sysplex environment, the data set names on TRCOUTnn DD statements must be unique throughout the sysplex. An ENQUEUE error results if the data set names are not unique. *z/OS MVS Diagnosis: Tools and Service Aids* has a discussion about creating external writer procedures.

An example of values to use in allocating IBMUSER.CT1 is:

```
DATA SET NAME: IBMUSER.CT1

VOLUME SERIAL      ===> USRPAK    (BLANK FOR AUTHORIZED DEFAULT
GENERIC UNIT       ===>           (GENERIC GROUP NAME OR UNIT
SPACE UNITS        ===> CYLS      (BLKS, TRKS, OR CYLS)
```

```
PRIMARY QUANTITY   ===> 3         (IN ABOVE UNITS)
SECONDARY QUANTITY ===> 5         (IN ABOVE UNITS)
DIRECTORY BLOCKS   ===> 0         (ZERO FOR SEQUENTIAL DATA SET)
RECORD FORMAT      ===> VB
RECORD LENGTH      ===> 23472
BLOCK SIZE         ===> 23476
EXPIRATION DATE    ===>           (YY/MM/DD, YYYY/MM/DD
```

To start a CT external writer, enter the following command at the operator console:

```
TRACE CT,WTRSTART=CTWTR
```

When you see the following message, the CT external writer is now ready:

```
ITT110I INITIALIZATION OF CTRACE WRITER CTWTR COMPLETE.
```

Step 2. Telling NFS to connect to a CT external writer and start sending records

To have the NFS server start sending trace records to the CT external writer, enter the following command at the operator console:

```
TRACE CT,ON,COMP=MVSNFS
```

In response to operator message *nn ITT006A SPECIFY OPERAND(S) FOR TRACE CT COMMAND reply with:

```
R nn,WTR=CTWTR,OPTIONS=(value(,value)),END
```

NFS will now start sending its trace records to the CT external writer, which in turn will start writing them to IBMUSER.CT1. A description of NFS trace option values (FFDC, Entry, Exit, etc.) is listed under “Starting component tracing for the z/OS NFS server” on page 161.

Using IPCS to view records from an external writer

Before you can use IPCS to look at trace records captured by a CT external writer, the remaining trace buffers must be queued to the external writer and the writer must be stopped in order to place an end-of-file record at the end of the data set.

To stop NFS tracing and flush the remaining buffer to the external writer issue the command:

```
TRACE CT,OFF,COMP=MVSNFS
```

NFS will issue a message GFSA937I to indicate the trace has been flushed. The writer may be stopped once the message has been displayed.

Finally, you can issue:

```
TRACE CT,WTRSTOP=CTWTR
```

which will stop the writer and make the data set available to IPCS. You can now use the IPCS CTRACE command CTRACE COMP(MVSNFS) FULL to look at the trace records, where the default data set has been set to IBMUSER.CT1.

When writing NFS trace records from more than one NFS server into separate component trace data sets, allocate each data set on a separate volume for improved performance. Use the IPCS COPYTRC command to merge records from multiple component trace data sets into one data set. Each component trace data set may represent one NFS server in the NFS multi-server and z/OS sysplex implementations. See *z/OS MVS IPCS Commands* for additional information.

Setting up a dump data set for abnormal ends

Normally, the Network File System ESTAE issues a SVC dump when failure occurs. However, if ESTAE is not able to do this, z/OS takes over and issues the appropriate dump you coded in your DD statement. This is an example of setting up a DUMP data set:

```
//SYSMDUMP  DD   DISP=SHR,DSN=MVSNFS.SYSMDUMP
```

File attributes of this DUMP data set should be set up like this:

```
Organization       PS
Record Format      FB
Record Length      4160
Block Size         4160
```

Searching the IBM database for APARs and PTFs

If your installation has access to the interactive online database program, IBMLink/ServiceLink, you can use IBMLink/ServiceLink to perform these tasks:

1. Search and browse for an existing APAR that is similar to your problem. Use the full set of keywords that is developed from the diagnostic procedures. Use only the keywords that are described in this book. Make sure that keywords are spelled exactly as they are described in this book.
2. If an APAR exists, search to see if a program temporary fix (PTF) is available.
3. If a PTF is available, order the PTF.
4. If an APAR does not exist, you can create an Electronic Technical Response (ETR) problem report to receive assistance from a z/OS NFS service representative. See "Contacting the IBM Support Center" for the type of information you will need.

Contacting the IBM Support Center

If your installation does not have access to IBMLink/ServiceLink, z/OS NFS also provides telephone support. Within the U.S.A. and Puerto Rico, call the following number and request assistance for the z/OS NFS feature by specifying the program number 5695DF121 and release level keyword.

```
IBM Support Center
1-800-237-5511
Monday through Friday, 8 a.m.-5 p.m. (excluding national holidays)
```

Outside of the U.S.A. and Puerto Rico, contact your local IBM representative.

When contacting IBM, be prepared to supply the following information:

- Your customer number
- Release level
- Current service level (from the APAR list)
- Keyword set or sets used to search in IBMLink/Service

You will be asked to describe the Network File System server and client machine environment. The IBM support representative might request the following relevant information.

- A minimum set of input commands on the client machine or z/OS operator console that reproduces the error.
- A copy of the minimum output from the client machine or z/OS operator console necessary to illustrate the failure.

- A copy of the z/OS NFS server log data set, the z/OS NFS client log data set, or both log data sets created by the input commands provided to recreate the error.
- Storage dump (if for an abnormal end).
- Linkedit map (if for abnormal end).
- Other supporting material, such as trace file printout from a network analyzer.
- For DOC SC26-7417 (*z/OS Network File System Guide and Reference*) failures, include the revision number and page(s) containing the error, and a description of the problem it caused.
- A copy of the attributes data set.
- A copy of the exports data set.
- A copy of the Network File System startup procedure.

Submitting Documentation on Tape: If the IBM service representative requests you submit documentation on tape, please write it to a standard label tape and include a hard copy of the DCB information for each data set along with the JCL used to create the tape.

Diagnostic aids

A description of first failure data capture, including SVC dump characteristics, dump contents, and errors and messages, is provided with the NFS as a major diagnostic aid.

First failure data capture

Network File System RAS characteristics are improved by the capture of diagnostic service data. Error records are written to SYS1.LOGREC and dumps are requested to SYS1.DUMPnn (these are in addition to the existing server trace). Component-specific information is provided in the SYS1.LOGREC entry and in the dump for the generation of RETAIN® search symptom strings.

Symptom data

Table 53 lists the component-specific symptom data placed in the System Diagnostic Work Area (SDWA).

Table 53. NFS symptom data

SDWA Field	Meaning	RETAIN Key
SDWAMODN	Active load module name	RIDS/
SDWACSCT	Active CSECT name	RIDS/
SDWAMDAT	Active CSECT assembly date	VALU/C
SDWAMVRS	Active CSECT service level	VALU/C
SDWAREXN	Recovery routine module name	RIDS/
SDWARRL	Recovery routine label name	FLDS/
SDWACID	Component identifier	PIDS/
SDWACIDB	Base component identifier	PIDS/
SDWASC	Active server function name	RIDS/

SVC dump

The Network File System SVC dumps have these characteristics:

Dump title: The dump title contains the component name, component identifier, release level, abend code, reason code, and the name of the ESTAE module requesting the dump. If available, the name of the failing module and the offset within the module are included.

Dump content: Table 54 shows the dump options and areas of storage that are included in the dump request:

Table 54. Dump content and storage areas

Dump Options	Storage Areas
SUMDUMP	Suspend summary dump
RGN	Server private area storage; programs and subpools
TRT	z/OS trace table
GRSQ	GRS ENQ control blocks
IO	I/O data areas
ALLPSA	All Prefixed Storage Areas
DFA	Data Facility Area
DFVT	Data Facility Vector Table
NFSSVT	Network File System Vector Table

Eye-catchers: Each CSECT within each server load module is identified by the CSECT name, the compile date, and the FMID of "HDZ11VS" or APAR level. Each function within a CSECT is identified by its variable length name.

Dump suppression

z/OS Dump Analysis and Elimination (DAE) is supported by providing sufficient information for DAE to uniquely identify the dump and by setting the VRADAE key in the Variable Recording Area (VRA) of the SDWA.

Data capture suppression

SYS1.LOGREC entries and SVC dumps are not requested for these abend codes:

X'0F3'
X'806'
X'A03'
X'x13'
X'x22' (except X'122')
X'x37'
X'x3E'

Invocation

A server ESTAE instance is entered whenever any server task ends abnormally. It is the ESTAE's responsibility to ensure that adequate and correct diagnostics are captured.

Errors and messages

Table 55 on page 238 illustrates the diagnostic errors and messages GFSA470I and GFSA471I.

Message GFSA470I is written if the SVC dump request fails. The message contains the error reason code from the SDUMP service.

Secondary ESTAE routines detect failures during the execution of the primary server ESTAE. If ESTAE processing is unable to complete, message GFSA471I is issued, and the server task is allowed to stop. The message contains the last abend code detected by the secondary ESTAE routines.

Table 55. Diagnostic errors and messages

Message	Explanation
GFSA470I	NETWORK FILE SYSTEM SVC DUMP REQUEST FAILED. REASON = hh
GFSA471I	NETWORK FILE SYSTEM ESTAE EXIT UNABLE TO COMPLETE PROCESSING. ABEND = X'xxx'

Debug trace data capture

The z/OS Network File System uses trace facilities to record debug trace diagnostic information when a problem requires additional diagnostic information beyond diagnostic messages. Assuming that debug trace diagnostics were not activated at the time of the original failure, the problem must be recreated a second time and the NFS debug trace facilities turned on to capture the diagnostic information. The server provides debug trace diagnostic information from two separate trace facilities. The z/OS Network File System client provides debug trace diagnostic information from one trace facility.

The z/OS NFS server records z/OS NFS debug trace diagnostic information to the z/OS component trace buffer or to the server log data sets. z/OS NFS error and informational messages are also recorded to the z/OS component trace buffer or to the server log data sets. When recording debug trace diagnostic information, IBM recommends that the z/OS component trace buffer be used to minimize the performance impact of collecting these extra diagnostics. To record z/OS NFS server debug trace diagnostic information in the z/OS component trace buffer, see “Using z/OS component tracing” on page 229. To record z/OS NFS server debug trace diagnostic information in the log data sets, see “z/OS NFS server debug trace capture” on page 381.

The z/OS NFS client records debug trace diagnostic information (including error and informational messages) to the client log data sets. To record z/OS NFS client debug trace diagnostic information in the log data sets, see “z/OS NFS client debug trace capture” on page 383.

Environmental checklist

This environmental checklist covers useful information that is recommended prior to the initializing the NFS server.

Dispatching Priority

Ensure that z/OS NFS has lower dispatching priority than TCP/IP. Also ensure that TCP/IP has a lower dispatching priority than VTAM®.

TCP/IP

Ensure that the MTU (Packet Size) is set to the lowest MTU when in a heterogeneous network. For example, if the network is comprised of:

```
Ethernet 802.3 (MTU=1492)
Ethernet Version 2 IEEE (MTU=1500)
token ring (MTU=2000)
FDDI (MTU=4000)
CTC (MTU=65527)
CLAW (MTU=4096)
```

In this example, the lowest MTU is set to 1492 to reduce packet fragmentation. The MTU setting is defined in the TCP/IP profile.

NFS

Verify that NFS is fully initialized.

- Portmapper is up.
- NFS has obtained port 2049 for NFS program such as *rpinfo -p <hostname>*.
- NFS has z/OS UNIX SEGMENT with UID=0 defined. The NFS server also needs OPERATIONS RACF authority.

For more information, see "Configuring the z/OS NFS server" on page 133 and "Configuring the z/OS NFS client" on page 131.

Useful Utilities

The following utilities are available on client machines help diagnosing simple network connection:

- *rpinfo -p <hostname>* to determine if portmapper (port 111), NFS (port 2049) are initialized with the appropriate port.
- *ping* to confirm that there is a live TCP/IP connection between client/sever machines.
- *traceroute <hostname>* to determine how packets are being routed from client to server.
- *iptrace* (AIX) or *snoop* (Solaris) are useful packet tracing utilities used to debug inbound and outbound packets between client and server. For example, during a mount request, using *iptrace* or *snoop* will show whether the client transmitted the mount request, the server has received the request, or the server is still processing the request.

Chapter 19. Network File System messages

This chapter lists messages from the NFS server, the NFS client, and the client operating system.

Table 56 lists the contents of this chapter.

Table 56. Messages - client operating system, NFS server, and NFS client

Section	Page
Server messages	241
Client messages	281
Messages from the client platform (AIX)	295

Server messages

This is a listing of the messages generated by the NFS server. Each message description gives an explanation and recommended actions where applicable. The system substitutes data for any part of a message shown here in *italics*.

Messages appear in the NFS server log data set in the same format as this example: **19:25:14 GFSA348I (I) GFSAMAIN ANMAI 02 NFS_INIT: z/OS NETWORK FILE SYSTEM SERVER (HDZ11VS) STARTED.**

Table 57 on page 241 shows the message format for the NFS server log data set:

Table 57. Message format for the NFS server log data set

19:25:14	The time stamp (*hours*:*minutes*:*seconds*)
GFSA	The component identifier for the NFS server
348	A unique message number
I	The message type: **A** Action; the user must perform a specific action. **E** Eventual action; the user must perform an action when time is available. **I** Informational; no user action is required.
(I)	The message level: E (error), W (attention), or I (informational). The system programmer can use the message level to determine which type of messages are shown by specifying log=error, log=warn, or log=info.
GFSAMAIN ANMAI 02 NFS_INIT:	Programming support information
z/OS NETWORK FILE SYSTEM SERVER (HDZ11VS) STARTED.	The message text

The messages are listed in numerical order (the time stamp, message level, and programming support information are not shown).

© Copyright IBM Corp. 1991, 2006

1. A value of *h_digits* is a hexadecimal number, and *d_digits* is a decimal number. A value of *text* or *dsname* is variable text (such as a data set name).
2. Messages GFSA300I through GFSA319I are intended for IBM support personnel when they are performing diagnosis. They do not indicate a problem with NFS but do provide statistics on NFS processing. As such, an extensive number of GFSA300I through GFSA319I messages may be issued.
3. For messages written to the console, the name of the start procedure is substituted for *procname*.

GFSA320I (*procname*) NETWORK FILE SYSTEM SERVER INITIALIZATION FAILED: *text*

Explanation: In the message text:

procname The name of the start procedure.

text The value of *text* can be any of these messages:

- `VIRTUAL STORAGE IS UNAVAILABLE`
- `MAPPING SIDE FILE NOT FOUND`
- `MAPPING SIDE FILE HAS INVALID SYNTAX OR FORMAT`
- `ERROR OPENING/READING MAPPING SIDE FILE`
- `SIDE FILE SPECIFIED BUT MAPPING IS DIS-ALLOWED BY INSTALLATION`
- `TASK IS NOT APF AUTHORIZED`
- `DFP LEVEL MUST BE DFSMS 1.2 OR HIGHER`
- `SERVER ALREADY STARTED`

System action: The NFS server startup ends.

Operator response: Notify the system programmer.

System programmer response: The system programmer response follows:

- If it is a virtual storage problem, increase the region size.
- If it is an APF-authorization problem, APF authorize all libraries in the STEPLIB DD statement.
- If the mapping side file is not found, make sure that the name specified in the attribute data set is correct and the file exists.
- If the mapping side file has invalid syntax, check the part GFSAPMAP in SYS1.NFSSAMP library for mapping side file rules.
- If an error occurs during the opening of the side file, check to make sure that the side file is not migrated and that it is readable.
- If **sfmax=0**, then the side file cannot be specified in the attribute data set.
- If you are using a down-level release of DFSMSdfp™, restart the NFS server after installing DFSMSdfp 1.2 or a later release.
- If the same start procedure is used with the same TCP/IP stack name, use a different name for either the start procedure or the TCP/IP stack name.

GFSA321I (*procname*) NETWORK FILE SYSTEM SERVER INITIALIZATION FAILED: z/OS UNIX ADDRESS SPACE HAS NOT BEEN STARTED.

Explanation: The Network File System server was not able to establish successful communications with the UNIX address space.

In the message text:

procname The name of the start procedure.

System action: The NFS server startup ends.

Operator response: Before starting the NFS server, both the UNIX and the TCP/IP address spaces must have been successfully started.

GFSA322I (*procname*) z/OS UNIX V_REG FAILED: RV=1, RC=*h_digit1*, RSN=*h_digit2*.

Explanation: The NFS server failed to register.

In the message text:

procname The name of the start procedure.

System action: The NFS server ends.

Operator response: Contact the system programmer.

System programmer response: The values of *h_digit1* and *h_digit2* are the return code and reason code from the z/OS UNIX V_REG callable service. See *z/OS UNIX System Services File System Interface Reference* for more information about return code and reason codes.

GFSA325I REQUESTED MEMORY NOT AVAILABLE.

Explanation: An operation to allocate virtual memory was tried but was unsuccessful. If this condition persists, the cause might be either of these problems.

- The value specified in the **region** parameter is too small.
- The value specified in the **bufhigh** attribute is too large.

System action: The request is stopped. The NFS server processing continues.

GFSA326E NETWORK FILE SYSTEM SERVER INITIALIZATION FAILED: CANNOT BIND *text1* PORT *d_digit* FOR PROGRAM *text2*.

Explanation: The NFS server was not able to bind port *d_digit*. The value of *text1* is the type of port. The value of *text2* is the program name.

System action: The NFS server ends.

Operator response: Notify your system programmer.

System programmer response: Program *text2* is either specified in /etc/services with a duplicate port *text1*, or is reserved in the range of the NFS server defined in TCPIP.PROFILE but port *text1* does not match program *text2* in /etc/services. Correct the situation and retry.

GFSA328I (*procname*) NO SWAP REQUEST FOR NETWORK FILE SYSTEM SERVER FAILED.

Explanation: The z/OS NO SWAP request for NFS failed.

In the message text:

procname The name of the start procedure.

System action: NFS server ends.

Operator response: Contact the system programmer.

System programmer response: Try to restart the NFS server. If the failure appears to be an NFS error, contact the IBM Support Center. Have available a symptom string and a copy of the z/OS console log.

GFSA329I (*procname*) SERVER SHUTDOWN IN PROGRESS.

Explanation: Shutdown procedures have started.

In the message text:

procname The name of the start procedure.

System action: The shutdown of the NFS server continues.

GFSA330I (*procname*) SERVER SHUTDOWN COMPLETE.

Explanation: The NFS server and its associated subtasks ends.

In the message text:

procname The name of the start procedure.

System action: NFS ends.

GFSA331E (*procname*) RECALL FAILED FOR MIGRATED DATA SET *dsname*.

Explanation: DFSMShsm was unable to recall a data set, *dsname*, because the data movement program, DFSMSdss, detected during a restore that the migrated data set had internal errors.

In the message text:

procname The name of the start procedure.

System action: NFS processing continues.

Operator response: Notify the system programmer for recovery actions.

System programmer response: The data set *dsname* had an internal error when it was migrated and cannot be recalled. See the DFSMShsm message ARC0075E for appropriate recovery actions.

GFSA333I (*procname*) z/OS NETWORK FILE SYSTEM SERVER *fmid* IS NOT STARTED BY A START COMMAND.

Explanation: The system programmer tried to start the NFS server with a command other than **start**.

In the message text:

procname The name of the start procedure.

fmid The FMID of the NFS server; for example, null or (HDZ11VS).

System action: NFS stops.

System programmer response: Start the NFS server by issuing a **start** command.

GFSA334E (*procname*) MOUNT HANDLE DATABASE CANNOT BE READ.

Explanation: The NFS server attempted to read the mount handle database but could not do so. This might occur because the mount handle database could not be opened or because it contained a record with an incorrect length. If migration coexistence APAR OA11612 is installed, pre-HDZ11TS mount handle databases are not supported.

In the message text:

procname The name of the start procedure.

System action: The NFS Server ends.

Operator response: Contact the system programmer.

System programmer response: If the mount handle database could not be opened, check that the record length of the mount handle database is 1600 for HDZ11TS and 2000 for HDZ11US/HDZ11VS. If the record length of the mount handle database is correct, then rebuild the mount handle database.

GFSA335E (*procname*) MOUNT HANDLE DATA SET CANNOT BE WRITTEN, EXPECTING LEN *d_digits1* REAL LEN *d_digits2*, VSAM R15(DEC) *d_digits3* REASON CODE(DEC) *d_digits4* LAST OP(DEC) *d_digits5*.

Explanation: The NFS server attempted to write a mount record to the mount handle database but could not do so. The length of the mount record is *d_digits1* bytes, but only *d_digits2* bytes were written. The failing information in writing the virtual storage access method (VSAM) key-sequenced data set (KSDS) mount handle database is the decimal return code *d_digits3*, the decimal error code or reason code *d_digits4*, and the code for the last operation *d_digits5*.

In the message text:

procname The name of the start procedure.

System action: System processing continues but in a degraded mode. Additional mounts or unmounts might fail in writing the record to the mount handle data set.

Operator response: Contact the system programmer.

System programmer response: The mount handle data sets have probably become unusable and need to be either cleared or restored to some previous level, and then the server has to be restarted. The information for the mount point was not saved in the mount handle data set. Clients might have to unmount and mount these mounted directories when the server is restarted.

GFSA335I MOUNT HANDLE DATA SET CANNOT BE WRITTEN, EXPECTING LEN *d_digits1* REAL LEN *d_digits2*, VSAM R15(DEC) *d_digits3* REASON CODE(DEC) *d_digits4* LAST OP(DEC) *d_digits5*.

Explanation: The NFS server tried to write a mount record to the mount handle database but could not do so. The length of the mount record is *d_digits1* bytes, but only *d_digits2* bytes were written. The failing information in writing the VSAM KSDS mount handle database is the decimal return code *d_digits3*, the decimal error code or reason code *d_digits4*, and the code for the last operation *d_digits5*.

System action: System processing continues but in a degraded mode. Additional mounts or unmounts might fail in writing the record to the mount handle data set.

Operator response: Contact the system programmer.

System programmer response: The mount handle data sets have probably become unusable and need to be either cleared or restored to some previous level, and then the server has to be restarted. Clients might have to unmount and mount any previously mounted directories.

GFSA336E (*procname*) MOUNT HANDLE DATA SET CANNOT BE OPENED, VSAM R15(DEC) *d_digits1* REASON CODE(DEC) *d_digits2* LAST OP(DEC) *d_digits3*.

Explanation: During a resource timeout, the NFS server tried to open the mount handle data set for writing but could not do so. The failing information in **fopen** for the VSAM KSDS mount handle data set is the decimal return code *d_digits1*, the decimal error code or reason code *d_digits2*, and the code for the last operation *d_digits3*.

In the message text:

procname The name of the start procedure.

System action: NFS stops.

Operator response: Contact the system programmer.

System programmer response: The mount handle data sets probably were created with incorrect attributes.

GFSA337E (*procname*) Network File System Server CTRACE initialization failed because macro *text1* had return code=*digit1*, reason code=*digit2*. Diagnostic information is being recorded in error log data sets only.

Explanation: The Network File System server was not able to obtain services needed for NFS component tracing. This message is received during the Network File System server initialization or resulting from a TRACE CT,ON,COMP=MVSNFS command.

In the message text:

procname
The name of the start procedure.

Text1 Name of system service call which failed.

Digit1 Return code returned from the service.

Digit2 Reason code returned from the service.

System action: The Network File System server continues. Debug trace messages, if any, will be recorded to the error log data sets only.

Operator response: Notify the system programmer.

System programmer response: Return and reason codes for system services such as DSPSERV, ALESERV and CTRACE are documented in *z/OS MVS Programming: Assembler Services Reference ABE-HSP* and *z/OS MVS Programming: Authorized Assembler Services Reference ALE-DYN*. Restart the NFS server to enable NFS component tracing after correcting the error or contact IBM support.

Source: z/OS NFS Server

GFSA338E (*procname*) UNKNOWN ENTERED OPTION=*text*

Explanation: Option *text* entered on the operator console for the z/OS NFS Server component trace command was invalid.

In the message text:

procname The name of the start procedure.

text Option text entered for NFS component trace on the operator console.

System action: None. The invalid option text is ignored.

Operator response: Correct the option and repeat the TRACE CT command to specify the necessary option. See the Operations chapter of*z/OS Network File System Guide and Reference* for a list of valid component trace options for the z/OS NFS server.

System programmer response: None.

GFSA339I (*procname*) Network File System *text1* RESTRICTION FOR DATA SET *text2*; ACCESS BLOCKED TO *text3* DATA SETS.

Explanation: NFS encountered *text2* data set that is not supported at NFS release level *text1*. *text3* can be the following depending on the *text1* release level:

- HDZ11TS:
  1. COMPRESSED
  2. 64K TRACK
  3. MULTI-VOLUME
- HDZ11US:
  1. COMPRESSED
  2. 64K TRACK

In the message text:

| **procname**
| The name of the start procedure.

| **Text1** NFS release level.

| **Text2** Name of unsupported data set.

| **Text3** Type of data set

System action: NFS processing terminates.

Operator response: Notify the system operator.

System programmer response: Remove the data set and retry NFS.

Source: z/OS NFS Server

GFSA340I (*procname*) Mvslogin is replacing userid *text1* with userid *text2* for client UID: *d_digits*, IP: *text3*

Explanation: During MVSLOGIN, the existing MVS userid, *text1*, having RACF authorization has been replaced with the new MVS userid, *text2*, having RACF authorization for the client with UNIX UID *d_digits* on the host system with IP address *text3*.

In the message text:

| **procname**
| The name of the start procedure.

| **Text1** The existing MVS userid.

| **Text2** The new MVS userid.

| **d_digits**
| UNIX user identification (UID).

| **Text3** IP address of the host system.

System action: NFS processing continues with the new MVS userid authority.

Operator response: None.

System programmer response: None.

Source: z/OS NFS Server

GFSA344I (*procname*) MOUNT HANDLE DATABASE CANNOT BE READ.

Explanation: The NFS server attempted to read the mount handle database but could not do so. This might occur because the mount handle database contained a record with an incorrect length or incorrect contents.

In the message text:

procname The name of the start procedure.

System action: System processing continues but in a degraded mode. Some or all the directories have not been remounted.

Operator response: Contact the system programmer.

System programmer response: If the mount handle database contains a record with an incorrect length or incorrect contents, then rebuild the mount handle data sets.

| **GFSA345E VSAM DATA SET IS NOT REUSABLE.**

| **Explanation:** At least one VSAM data set (FHDBASE, | FHDBASE2, LDBASE, or LDBASE2 DD statements in | the z/OS NFS Server job) is not defined with the | REUSE option.

| **System action:** NFS server processing stops.

| **Operator response:** Contact the system programmer.

| **System programmer response:** Recreate the | appropriate VSAM data sets using IDCAMS with the

REUSE option, and restart the NFS server.

GFSA346I *time_stamp*.

Explanation: Displays the current time stamp. This message is issued when the NFSLOG switches.

GFSA347I ERROR RETURNED TO CLIENT: RC = *d_digits text*.

Explanation: The error code *d_digits* was returned to the client.

In the message text:

text The value of *text* is the meaning of the error code.

System action: NFS processing continues.

GFSA348I (*procname***) z/OS NETWORK FILE SYSTEM SERVER** *fmid* **STARTED.**

Explanation: The NFS is initialized and ready to accept **modify** commands from the operator console.

In the message text:

procname The name of the start procedure.

fmid The FMID of the NFS server; for example, null or (HDZ11VS).

System action: NFS continues processing.

GFSA349I UNEXPECTED ERROR DETECTED: *d_digits text*.

Explanation: The NFS server encountered a condition that indicates continued processing might produce undesirable results.

In the message text:

text The value of *text* is additional debugging information for programming support personnel.

System action: NFS either shuts down or stops the request and continues processing, depending on where the error was detected.

System programmer response: Contact your programming support personnel.

GFSA350E Lock Data Set ddname cannot be opened, VSAM rc=d_digits1 rsnc=d_digits2 LastOp=d_digits3

Explanation: he NFS server tried to open a locking data set but could not do so. This message identifies the data set, the return and reason codes, and the code of the last operation.

In the message text:

ddname
The ddname of the failing data set.

d_digits1
The failing information in **fopen** for the VSAM KSDS data set -- the decimal return code.

d_digits2
The decimal error code or reason code.

d_digits3
The code for the last operation.

System action: If this is the first of the two lock data sets to fail, NFS continues using the other lock data set. If this is the second of the two lock data sets to fail, lock information will not be recorded and reclaim permissions on the next NFS restart may be inaccurate.

Operator response: Notify the system programmer.

System programmer response: Verify that the lock data set is allocated because it must always be allocated, even if *nonlm* is specified in the site attributes. If the lock data set has been allocated, then it has probably been corrupted. Plan to delete the data set and reallocate it. You may also wish to verify that the ddname in the NFS startup job refers to the same data set as was created for it.

Source: z/OS NFS Server

GFSA351E Lock Data Set *ddname* **(***dsname***) is not proper VSAM format: VSAM=***d_digits1* **type=***d_digits2* **(expect** *d_digits3***) KeyOff=***d_digits4* **(expect 0) KeySize=***d_digits5* **(expect** *d_digits6***) RecSize=***d_digits7* **(expect** *d_digits8***)**

Explanation: The Network File System server detected that the lock data set does not have the proper characteristics for a VSAM key sequenced data set (KSDS). This message identifies the incorrect characteristics and their expected values.

In the message text:

ddname
The ddname of the failing data set.

dsname
The data set name of the failing data set.

d_digits1
Indicates VSAM data set or not: value of 1 indicates VSAM, value of 0 indicates not VSAM.

d_digits2
The detected type of the data set.

d_digits3
The expected type of the data set (KSDS).

d_digits4
The detected offset of the VSAM key within the record, which was expected to be 0.

d_digits5
The detected size of the VSAM key within the record.

d_digits6
The expected size of the VSAM key within the record.

d_digits7
The detected size of the VSAM record.

d_digits8
The expected size of the VSAM record.

System action: This lock data set will not be used. NFS will continue using the other locking data set. If both lock data sets have failed, lock information will not be recorded and reclaim permissions on the next NFS restart may be inaccurate.

Operator response: Notify the system programmer.

System programmer response: The lock data set has been created incorrectly. Reallocate the lock data set with correct parameters, as described in "Lock data sets" on page 139.

Source: z/OS NFS Server

GFSA352E Lock Data Set ddname failed op opname with rc=d_digits1

Explanation: An attempt to read, write or modify the lock data set has failed. This messages identifies the data set, the operation and the failure code.

In the message text:

ddname
The ddname of the failing data set.

opname
The name of the operation that failed.

d_digits1
The failure code from the operation.

System action: This lock data set will not be used. NFS continues, using the other lock data set. If both lock data sets have failed, lock information will not be recorded and reclaim permissions on the next NFS restart may be inaccurate.

Operator response: Notify the system programmer.

System programmer response: The lock data set has probably been corrupted. Plan to delete the data set and reallocate it.

Source: z/OS NFS Server

GFSA360I *text*.

Explanation: This message displays memory management statistics, *text*.

System action: NFS continues processing.

GFSA361I NETWORK FILE SYSTEM IS SHORT ON STORAGE.

Explanation: This message displays on the operator console when a shortage of virtual storage is detected.

System action: NFS processing continues. The storage constraint might be relieved when some storage is freed up later on.

Operator response: If this message is displayed repeatedly within a short period of time, stop or cancel NFS and notify the system programmer.

System programmer response: Take either or both of the following actions before restarting the NFS server.

- Increase the region size for the step or started task.
- Decrease the value specified for the **bufhigh** attribute of the attributes data set.

GFSA362I (*procname*) REGION SIZE WILL NOT ACCOMMODATE BUFHIGH AND LOGICAL CACHE SPECIFICATIONS.

Explanation: The values of **bufhigh** and **logicalcache** specified in the attributes data set are incorrect.

In the message text:

procname The name of the start procedure.

System action: NFS ends.

System programmer response: Either increase the region size of the step, or reduce the total value of the **bufhigh** and **logicalcache** attribute values.

GFSA363I (*procname*) NETWORK FILE SYSTEM SERVER IS SHORT OF BUFFERS.

Explanation: NFS has no memory blocks in the buffer area limited by the **bufhigh** attribute for data buffers.

In the message text:

procname The name of the start procedure.

System action: None

Operator response: Notify the system programmer.

System programmer response: Increase the **bufhigh** value in the NFS attribute profile.

GFSA364I (*procname*) Not enough storage below 16Mb for requested number of Legacy Tasks.

Explanation: The size of available storage is not enough for the requested number of legacy Tasks.

In the message text:

procname The name of the start procedure.

System action: The NFS server will be shut down.

User response: None

Operator response: Contact the system programmer.

System programmer response: Reduce the number of legacy tasks in the site attribute data set, using the **nfstasks** attribute.

GFSA400I INVALID RECFM SPECIFICATION (*text*).

Explanation: The value of *text* is the incorrect record format that was specified in the attributes data set.

System action: NFS stops.

System programmer response: Correct the attributes data set.

GFSA401I (*procname*) CANNOT OPEN THE ATTRIBUTE DATA SET.

Explanation: The server was unable to open the attributes data set defined in the JCL for DDNAME NFSATTR. The DD statement might be missing or the data set name might be incorrect.

In the message text:

procname The name of the start procedure.

System action: NFS stops.

Operator response: Notify the system programmer.

System programmer response: Correct the JCL for DDNAME NFSATTR.

GFSA402I (*procname*) READ FAILED FOR THE ATTRIBUTES DATA SET.

Explanation: An error occurred while NFS was processing the attribute data set. This message follows other messages that describe the error in greater detail. The attributes data set is defined in the JCL for DDNAME NFSATTR.

In the message text:

procname The name of the start procedure.

System action: NFS stops.

Operator response: Notify the system programmer.

System programmer response: Correct the attributes data set.

Note: When the NFS attribute data set is created, the *num off* option in ISPF should also be used.

The sequence number is not allowed in the NFS attributes data set.

GFSA403I (*procname*) PARSE FAILED IN LINE *d_digits text*.

Explanation: The parsing of line number *d_digits* in the attribute data set failed.

In the message text:

procname The name of the start procedure.

text The value of *text* is the actual line from the attributes data set that contains the failure.

This message follows other messages that describe the error in greater detail.

System action: NFS stops.

System programmer response: Correct the attributes data set.

GFSA404I UNEXPECTED END OF STRING ON END OF PARSE IN LINE *d_digits*.

Explanation: A comma is missing between attributes on line number *d_digits* of the attributes data set.

System action: NFS stops.

System programmer response: Correct the attributes data set.

GFSA405I PARSE FAILED FOR ATTRIBUTE FIELD - ILLEGAL KEYWORD IN LINE *d_digits*.

Explanation: The keyword specified in line number *d_digits* of the attribute data set is not a valid attribute keyword.

System action: NFS stops.

System programmer response: Correct the attributes data set.

GFSA406I MISSING LEFT PARENTHESIS IN LINE *d_digits*.

Explanation: An attribute specified on line number *d_digits* of the attributes data set is missing a left parenthesis.

System action: NFS stops.

System programmer response: Correct the attributes data set.

GFSA407I MISSING RIGHT PARENTHESIS IN LINE *d_digits*.

Explanation: An attribute specified on line number *d_digits* of the attributes data set is missing a right parenthesis.

System action: NFS stops.

System programmer response: Correct the attributes data set.

GFSA408I PARSE FAILED ON NUMBER FIELD IN LINE *d_digits*.

Explanation: An attribute with a negative number was specified on line number *d_digits* of the attributes data set.

System action: NFS stops.

System programmer response: Correct the attributes data set.

GFSA409I INVALID DSORG SPECIFICATION IN LINE *d_digits*.

Explanation: The data set organization specified in the **dsorg** attribute on line number *d_digits* of the attributes data set is incorrect or is not supported by NFS.

System action: NFS stops.

System programmer response: Correct the attributes data set.

GFSA411I INVALID VOL SPECIFICATION IN LINE *d_digits*.

Explanation: The volume specified in the **vol** (volume) attribute on line number *d_digits* of the attributes data set is incorrect.

System action: NFS stops.

System programmer response: Correct the attributes data set.

GFSA412I INVALID UNIT SPECIFICATION IN LINE *d_digits*.

Explanation: The unit specified in the unit attribute on line number *d_digits* of the attributes data set is incorrect.

System action: NFS stops.

System programmer response: Correct the attributes data set.

GFSA413I THE ATTRIBUTE VALUE *d_digits1* IS NOT IN THE RANGE OF *d_digits2* TO *d_digits3*.

Explanation: The value of *d_digits1*, specified in one of the attributes of the attributes data set, must be between the minimum value, *d_digits2*, and the maximum value, *d_digits3*, for this attribute. Message GFSA403I follows this message.

System action: NFS stops.

System programmer response: See message GFSA403I to determine the attribute in error, and then correct the attributes data set.

GFSA414I THE ATTRIBUTE VALUE *d_digits1* EXCEEDS THE MAXIMUM TIMEOUT VALUE OF *d_digits2*.

Explanation: In the message text:

d_digits1 The value of *d_digits1*, specified in one of the attributes of the attributes data set, must be less than or equal to *d_digits2*.

d_digits2 The value of *d_digits2* is the maximum value allowed for the attribute.

Message GFSA403I follows this message.

System action: NFS stops.

System programmer response: See message GFSA403I to determine the attribute in error, and then correct the attributes data set.

GFSA415I THE ATTRIBUTE TIME OUT VALUE *d_digits1* IS LESS THAN THE MINIMUM TIME OUT VALUE *d_digits2*.

Explanation: The value of *d_digits1*, specified in the **attrtimeout** attribute in the attributes data set, must be greater than or equal to the value *d_digits2*, which is specified in the **mintimeout** attribute of the attributes data set.

System action: NFS stops.

System programmer response: Correct the attributes data set.

GFSA416I THE READ TIME OUT VALUE *d_digits1* IS LESS THAN THE MINIMUM TIME OUT VALUE *d_digits2*.

Explanation: The value of *d_digits1*, specified in the **readtimeout** attribute of the attributes data set, must be greater than or equal to the value *d_digits2*, which was specified in the **mintimeout** attribute of the attributes data set.

System action: NFS stops.

System programmer response: Correct the attributes data set.

GFSA417I THE WRITE TIME OUT VALUE *d_digits1* IS LESS THAN THE MINIMUM TIME OUT VALUE *d_digits2*.

Explanation: The value of *d_digits1*, specified in the **writetimeout** attribute of the attributes data set, must be greater than or equal to the value *d_digits2*, which was specified in the **mintimeout** attribute of the attributes data set.

System action: NFS stops.

System programmer response: Correct the attributes data set.

GFSA418I THE ATTRIBUTE TIME OUT VALUE *d_digits1* **IS GREATER THAN THE MAXIMUM TIME OUT VALUE** *d_digits2***.**

Explanation: The value of *d_digits1*, specified in the **attrtimeout** attribute of the attributes data set, must be less than or equal to the value *d_digits2*, which was specified in the **maxtimeout** attribute of the attributes data set.

System action: NFS stops.

System programmer response: Correct the attributes data set.

GFSA419I THE READ TIME OUT VALUE *d_digits1* **IS GREATER THAN THE MAXIMUM TIME OUT VALUE** *d_digits2***.**

Explanation: The value of *d_digits1*, specified in the **readtimeout** attribute of the attributes data set, must be less than or equal to the value *d_digits2*, which was specified in the **maxtimeout** attribute of the attributes data set.

System action: NFS stops.

System programmer response: Correct the attributes data set.

GFSA420I THE WRITE TIME OUT VALUE *d_digits1* **IS GREATER THAN THE MAXIMUM TIME OUT VALUE** *d_digits2***.**

Explanation: The value of *d_digits1*, specified in the **writetimeout** attribute of the attributes data set, must be less than or equal to the value *d_digits2*, which was specified in the **maxtimeout** attribute of the attributes data set.

System action: NFS stops.

System programmer response: Correct the attributes data set.

GFSA421I THE NOATTRTIMEOUT ATTRIBUTE WAS SPECIFIED BUT THE MAXTIMEOUT VALUE WAS SET TO *d_digits***.**

Explanation: The **noattrtimeout** attribute means that the data set is not to be deallocated following a **lookup** or **getattr** operation. The **maxtimeout** attribute specifies the maximum timeout value allowed for any of the timeout attributes. The value of *d_digits* was specified as the **maxtimeout** value, in seconds, that the data set is to remain allocated. These attributes are in conflict. The **noattrtimeout** and **maxtimeout** attributes are specified in the attributes data set.

System action: NFS stops.

System programmer response: If you want to use the **noattrtimeout** attribute, specify the **nomaxtimeout** attribute in the attributes data set. Correct the attributes data set.

GFSA422I THE NOREADTIMEOUT ATTRIBUTE WAS SPECIFIED BUT THE MAXTIMEOUT VALUE WAS SET TO *d_digits***.**

Explanation: The **noreadtimeout** attribute means that the data set is not to be deallocated following a read operation. The **maxtimeout** attribute specifies the maximum timeout value allowed for any of the timeout attributes. The value of *d_digits* was specified as the value of the **maxtimeout** attribute, in seconds, that the data set is to remain allocated. These attributes are in conflict. The **noreadtimeout** and **maxtimeout** attributes are specified in the attributes data set.

System action: NFS stops.

System programmer response: If you want to use the **noreadtimeout** attribute, specify the **nomaxtimeout** attribute in the attributes data set. Correct the attributes data set.

GFSA423I THE NOWRITETIMEOUT ATTRIBUTE WAS SPECIFIED BUT THE MAXTIMEOUT VALUE WAS SET TO *d_digits***.**

Explanation: The **nowritetimeout** attribute means that the data set is not to be deallocated following a write operation. The **maxtimeout** attribute specifies the maximum timeout value allowed for any of the timeout attributes. The value of *d_digits* was specified as the value of the **maxtimeout** attribute, in seconds, that the data set is to remain allocated. These attributes are in conflict. The **nowritetimeout** and **maxtimeout** attributes are specified in the attributes data set.

System action: NFS stops.

System programmer response: If you want to use the **nowritetimeout** attribute, specify the **nomaxtimeout** attribute in the attributes data set. Correct the attributes data set.

GFSA424I MINIMUM TIME OUT VALUE, *d_digits1***, IS GREATER THAN THE MAXIMUM TIME OUT VALUE,** *d_digits2***.**

Explanation: The value of *d_digits1*, specified in the **mintimeout** attribute of the attributes data set is greater than the value *d_digits2*, specified in the **maxtimeout** attribute of the attributes data set.

System action: NFS stops.

System programmer response: Correct the attributes data set.

GFSA425I INVALID SPECIFICATION RECFM(*text*).

Explanation: One of the characters in the value of *text* specified in the **recfm** attribute of the attributes data set is incorrect.

System action: The NFS stops.

System programmer response: Correct the attributes data set.

GFSA426I INVALID RECFM(*text*) - MUST SPECIFY U, F, OR V.

Explanation: One of the characters in the value of *text* specified in the **recfm** attribute of the attributes data set must define whether the record is fixed length (F), variable length (V), or undefined (U).

System action: NFS stops.

System programmer response: Correct the attributes data set.

GFSA429I INVALID DSNTYPE SPECIFICATION IN LINE *d_digits*.

Explanation: The data set type specified in the **dsntype** attribute on line number *d_digits* of the attributes data set is incorrect.

System action: NFS stops.

System programmer response: Correct the attributes data set.

GFSA430I INVALID *SMS_keyword* SPECIFICATION IN LINE *d_digits*.

Explanation: The SMS keyword *SMS_keyword* is syntactically incorrect on line number *d_digits*. See system-managed storage documentation for DATACLAS, MGMTCLAS, and STORCLAS naming conventions.

System action: NFS startup ends if the keyword was specified as a site attribute. If the incorrect SMS keyword was specified by a client as a mount parameter or in a command, the line number is set to zero and an I/O error is returned to the client.

System programmer response: Correct the site attributes file, if applicable.

GFSA431I INVALID OPTION SPECIFICATION OF *text* IN LINE *d_digits*.

Explanation: The option specified in the value of **text** on line number *d_digits_* of the attributes data set is incorrect.

System action: NFS stops.

System programmer response: Correct the site attributes file, if applicable.

GFSA434I (*procname*) *text1*(*d_digits1*) IS SET TO THE DEFAULT VALUE, *text1*(*d_digits2*,*d_digits3*,*d_digits4*).

Explanation: The value of *text1*(*d_digits1*) is the value from the previous release. This value could not be applicable to the new release. For forward compatibility, this value takes on the default value for the new release.

In the message text:

procname The name of the start procedure.

System action: NFS processing continues.

GFSA435E SUM OF *d_digit1* PLUS *d_digits2* MUST BE LESS THAN OR EQUAL TO *d_digit3*.

Explanation: The sum of *d_digit1* plus *d_digit2* is greater than *d_digit3*.

System action: NFS processing stops.

System programmer response: Correct the values.

GFSA436I INVALID SIDEFILE SPECIFICATION IN LINE *d_digits*

Explanation: A data set name was not specified in the **sidefile** attribute.

System action: NFS server startup ends.

System programmer response: Correct the problem and make the necessary changes in the attributes data set.

GFSA437I INVALID PUBLIC SPECIFICATION IN LINE *d_digits*

Explanation: Parsing of the **public** keyword resulted in an error for one of the following reasons:

- **public** keyword specification is syntactically incorrect
- No public path names have been specified
- The path name specified is not valid
- The HFS public path name does not match the HFS prefix

System action: NFS server startup ends.

Operator response: Correct the problem and make the

necessary changes in the attributes data set.

GFSA438I EXPORT SPANNING PATHNAMES NOT SUPPORTED

Explanation: The export-spanning path names support for a multicomponent **lookup** request is not supported.

System action: The request fails. NFS processing continues.

User response: Construct a different path name in which the path is not spanned.

GFSA439I z/OS UNIX PUBLIC PATHNAME SPECIFIED BUT z/OS UNIX IS NOT ENABLED

Explanation: A z/OS UNIX public path name was specified for the **public** keyword but **nohfs** was also specified, which disables z/OS UNIX processing.

System action: NFS server startup ends.

Operator response: Correct the problem and make the necessary changes in the installation table.

GFSA440I INVALID SECURITY SPECIFICATION IN *d_digits*

Explanation: Parsing of the **security** keyword resulted in an error for one of the following reasons.

- Missing first parameter
- First parameter not valid

System action: NFS server startup ends.

Operator response: Correct the problem and make the necessary changes in the attributes data set.

GFSA441E (*procname***) Invalid leasetime value. It must be smaller than logout value.**

Explanation: The specified Network File System lock lease time value is larger than the logout timeout value. The leasetime value must be smaller that the logout timeout value.

System action: The Network File System server initialization fails and the server terminates.

Operator response: None.

System programmer response: Check and correct the relationship between the leasetime and logout values specified in the Site Attribute File.

Source: z/OS NFS Server

GFSA450I CREATED TASK(*h_digits***) -** *text1* **-** *text2***.**

Explanation: The NFS is creating the number of tasks requested in the **nfstasks** attribute of the attributes data set. This message displays for each task created.

In the message text:

h_digits The value of *h_digits* is the TCB address.

text1 The value of *text1* is the task name.

text2 The value of *text2* is the module name.

System action: NFS processing continues.

GFSA451I DELETING TASK(*h_digits***) -** *text***.**

Explanation: The NFS is deleting a task. This is in response to the **stop** operand of the **modify** command. This message is displayed for each task deleted.

In the message text:

h_digits The value of *h_digits* is the TCB address.

text The value of *text* is the module name.

System action: NFS shutdown continues.

GFSA452I SUBTASK TERMINATED: *h_digits***.**

Explanation: NFS is stopping a task. This is in response to the **stop** operand of the **modify** command. This message is displayed for each ended task.

In the message text:

h_digits The value of *h_digits* is the task control block (TCB) address.

System action: NFS shutdown continues.

GFSA453E (*procname***) NETWORK FILE SYSTEM SERVER LOST TASK** *tasktype***. NOW AVAILABLE:** *nn1* **LEGACY TASKS,** *mm1* **HFS TASKS,** *oo1* **LONG SERVICE TASKS. LEGACY TASKS SHUTDOWN LIMIT IS <** *nn2***. HFS TASKS SHUTDOWN LIMIT IS <***mm2***. LONG SERVICE TASKS SHUTDOWN LIMIT IS <** *oo2***.**

Explanation: This message means that one *tasktype* task was detached due to an ABEND 878 . If the count of the available tasks of one type is less than the shutdown limit for the task, the NFS server restarts the suspended tasks.

In the message text:

procname The name of the start procedure.

tasktype This value indicates the type of tasks: "`legacy`", "`HFS`", or "`Long service`" tasks.

nn1 This value indicates the number of available legacy tasks (for conventional MVS data sets).

mm1 This value indicates the number of HFS tasks.

oo1 This value indicates the number of long service tasks.

nn2 This value indicates the limit of legacy tasks (for conventional MVS data sets).

mm2 This value indicates the limit of HFS tasks.

oo2 This value indicates the limit of long service tasks.

System action: The task ends. The NFS server restarts the suspended tasks.

User response: None

GFSA454I (*procname***)** *taskname* **TASK HAS BEEN RECOVERED AFTER 878 ABEND. ALL TASKS WILL BE STARTED NOW**

Explanation: The NFS task had an ABEND 878 (80A) and then was recovered. The NFS server is ready to restart task processing.

In the message text:

procname The name of the start procedure.

taskname Specifies the name of the NFS task that had the ABEND 878 (80A).

System action: NFS server restarts task processing.

User response: None

GFSA455E (*procname***)Connection failed from TCPIP, shutting down. Errno(***digit1***) errno2(***digit2***).**

Explanation: The Network File System was unable to connect to TCPIP The attempted connection returned error code *digit1* and reason code *digit2*.

System action: The Network File System server initialization fails and the server terminates.

Operator response: None.

System programmer response: Check the return and reason codes. If these indicate a system configuration error, correct it. Otherwise, submit this error to IBM Support.

Source: z/OS NFS Server

GFSA456E (*procname***) Socket creation failed, errno(***digit1***) errno2(***digit2***).**

Explanation: The Network File System was unable to create a communication socket. The attempted connection returned error code *digit1* and reason code *digit2*.

System action: The Network File System server initialization fails and the server terminates.

Operator response: None.

System programmer response: Check the return and reason codes. If these indicate a system configuration error, correct it. Otherwise, submit this error to IBM Support.

Source: z/OS NFS Server

GFSA459E (*procname***) ioctl: cannot set socket to blocking/non-blocking mode, errno(***digit1***) errno2(***digit2***).**

Explanation: The Network File System was unable to set the required socket blocking mode. The z/OS UNIX ioctl function returned error code *digit1* and reason code *digit2*.See *z/OS XL C/C++ Run-Time Library Reference* for details on the ioctl function.).

System action: The request require this mode fails.

Operator response: None.

System programmer response: Check the return and reason codes. If these indicate a system configuration error, correct it. Otherwise, submit this error to IBM Support.

Source: z/OS NFS Server

GFSA460E (*procname***) send/sendto request failed, errno(***digit1***) errno2(***digit2***).**

Explanation: The Network File System server was unable to send a request to unregister functions in the portmapper before registering this new instance of the server. z/OS UNIX returned error code *digit1* and reason code *digit2*.

System action: The Network File System server initialization fails and the server terminates.

Operator response: None.

System programmer response: Check the return and reason codes. If these indicate a system configuration error, correct it. Otherwise, submit this error to IBM Support.

Source: z/OS NFS Server

GFSA461E (*procname***) Min** *d_digits* **subtask threshold is reached.**

Explanation: The number of NFS Server subtasks has reached the minimum value.

In the message text:

procname The name of the start procedure.

System action: NFS stops.

User response: None

Operator response: Contact the system programmer.

System programmer response: Check mintasks and nfstasks in the NFS server attribute data set, and reduce this number if possible.

GFSA470I (*procname*) NETWORK FILE SYSTEM SERVER SVC DUMP REQUEST FAILED. REASON=*reason_code*.

Explanation: A request to write a z/OS SVC dump failed.

In the message text:

procname The name of the start procedure.

reason_code The value of *reason_code* is a hexadecimal number indicating the reason that z/OS was unable to write the dump. See the description of the **sdump** macro in *z/OS MVS Programming: Authorized Assembler Services Reference LLA-SDU* for the meaning of the reason code.

System action: ABEND processing continues.

Operator response: If the dump could not be written due to an operational procedure and a dump is necessary to diagnose the failure, correct the procedure.

System programmer response: If the failure appears to be due to an NFS error, contact the IBM Support Center. Have available a symptom string and a copy of the z/OS console log.

GFSA471I (*procname*) NETWORK FILE SYSTEM SERVER ESTAE EXIT UNABLE TO COMPLETE PROCESSING. ABEND=*abend_code*.

Explanation: The NFS ESTAE exit routine that ended abnormally and is unable to complete processing.

In the message text:

procname The name of the start procedure.

abend_code The *abend_code* value is set to the last abend code encountered by a secondary instance of the ESTAE exit routine.

System action: ESTAE processing is stopped. The NFS address space might end also.

Operator response: Restart the NFS address space. Notify network users of the failure.

System programmer response: Contact the IBM Support Center. Have available a symptom string, the related SYS1.LOGREC entries, any related SDUMPs, and the NFS error trace log.

GFSA472I (*procname*) 878 ABEND HANDLING FOR TASK *taskname* IN PROCESS. ALL TASKS STOPPED. PLEASE WAIT FOR RECOVERY.

Explanation: The NFS task had an ABEND 878 (80A), and the recovery procedure is in process.

In the message text:

procname The name of the start procedure.

taskname Specifies the name of the NFS task that had the ABEND 878 (80A).

System action: The NFS server is in recovery processing and all normal processing is suspended.

User response: None

GFSA473I (*procname*) No connection yet between *string1* UNIX & TCPIP. Waiting for connection. errno(*digit1*) errno2(*digit2*).

Explanation: The Network File System was unable to connect to TCP/IP. The attempted connection returned error code *digit1* and reason code *digit2*.

In the message text:

procname The name of the start procedure.

System action: NFS processing waits for TCP/IP connection to be established.

User response: None

Operator response: Contact the system programmer.

System programmer response: Check the error and reason codes. If these indicate a TCP/IP system configuration error, correct it. Otherwise, submit this error to IBM TCP/IP Support.

GFSA480I REQUEST FAILED BECAUSE OF ERRORS ENCOUNTERED IN THE CONVERSION SERVICES, RC=*d_digits1*, RSN=*d_digits2*.

Explanation: The request received an error condition from the conversion services while data was being translated.

System action: The request failed. NFS processing continues.

User response: See *z/OS Support for Unicode: Unicode Services* for information about the error encountered, and correct the problem if possible.

GFSA481I REQUEST FAILED BECAUSE TEXT CONVERSION RESULTED IN THE LENGTHS OF INPUT AND OUTPUT STRINGS BEING DIFFERENT.

Explanation: The request failed because the translation of the text string resulted in a different length.

Translations that involve length changes (such as SBCS to MBCS) are not currently supported.

System action: The request failed. NFS processing continues.

User response: Make the necessary changes so that text string translations do not result in different lengths.

GFSA482I REQUEST FAILED BECAUSE TRANSLATION OF TEXT IS NOT POSSIBLE.

Explanation: An operation is being requested for a file that is tagged with a valid coded character set identifier (CCSID), but the user has not specified a **cln_ccsid** value to be used for translation, and the **xlat()** keyword is also specified.

System action: The request failed. NFS processing continues.

User response: Either specify a **cln_ccsid** value for translation or remove the specification **xlat()** keyword.

GFSA483I WRITE REQUEST FAILED BECAUSE BINARY DATA CANNOT BE WRITTEN TO A FILE WITH PURE TEXT DATA.

Explanation: A write request has the **binary** option specified explicitly, and the file being written to has pure text data.

System action: The request failed. NFS processing continues.

User response: Either remove the **binary** option or change the file tag that is associated with the file to show that the file has mixed data.

GFSA483E TAG OPTION IS SPECIFIED BUT CONVERSION SERVICES IS NOT ACTIVE.

Explanation: Conversion Services must be active if TAG is specified in the site attribute data set or mount command.

System action: Network File System Server startup ends if TAG is specified in the site attribute dataset, or mount point is ignored during MHDB processing, or mount command is rejected.

Operator response: Notify the system programmer.

System programmer response: Either change from TAG to NOTAG in site attribute dataset (or mount command) or activate Conversion Services.

GFSA484E TECHNIQUE VALUE OF CONSERV OPTION IS INVALID.

Explanation: The technique specified for CONVSERV(technique) in the site attribute data set or mount command is incorrect. It can consist of up to five characters corresponding to the five available techniques: R, E, C, L and M.

System action: Network File System Server startup ends if CONVSERV(technique) is specified in the site attribute data set, or the mount command is rejected.

Operator response: Notify the system programmer.

System programmer response: Update the values for CONVSERV(technique).

GFSA485I (*procname*) CONVERSION SERVICES ARE NOT ACTIVE DURING NETWORK FILE SYSTEM SERVER STARTUP.

Explanation: During startup, the NFS Server detected that Conversion Services were not available.

In the message text:

procname The name of the start procedure.

System action: NFS server processing continues without the availability of Conversion Services.

Operator response: Notify the system programmer.

System programmer response: If data conversion is required, Conversion Services must be available to NFS. See *z/OS Support for Unicode: Unicode Services*for more information.

GFSA486I (*procname*) CONVERSION SERVICES ARE ACTIVE AND WILL BE USED BY NETWORK FILE SYSTEM SERVER.

Explanation: During RPC MOUNT processing, the NFS Server detected that Conversion Services were available.

In the message text:

procname The name of the start procedure.

System action: NFS server processing continues.

GFSA487E (*procname*) During startup, Network File System Server encountered problems with Conversion Services: Service returned RC = d_digit1, RSN = d_digit2.

Explanation: During Network File System server startup, it encountered problems with the Conversion Services. The conversion service failed.

In the message text:

procname The name of the start procedure.

d_digit1 The return code from conversion services.

d_digit2 The reason code from conversion services.

System action: NFS server processing continues.

System programmer response: The conversion services problem must be resolved before the conversion services can be used by the Network file system server.

GFSA501I REQUEST HEADER ALLOCATION FAILED.

Explanation: An operation to allocate virtual memory for a request header was tried but was unsuccessful.

System action: The request is stopped. NFS processing continues.

System programmer response: Increase the size of the step region.

GFSA550I (*procname*) IP*text* is active.

Explanation: The Network File System server found the IP level (*text*) of V6 or V4, with the TCPIP Stack active. The server will use that stack for all its communications.

System action: The Network File System server continues.

Operator response: None.

System programmer response: None.

Source: z/OS NFS Server

GFSA502I REQUEST HEADER DATA BLOCK ALLOCATION FAILED.

Explanation: An operation to allocate virtual memory for a request header data block was tried but was unsuccessful.

System action: The request is stopped. NFS processing continues.

System programmer response: Increase the size of the step region.

GFSA554I (*procname*) REGISTER FOR PORT *d_digits1* PROGRAM *d_digits2* - VERSION *d_digits 3* - FAILED.

Explanation: In the message text:

procname The name of the start procedure.

d_digits1 The value of *d_digits1* is the port number.

d_digits2 The value of *d_digits2* is the remote procedure call (RPC) program number.

d_digits3 The value of *d_digits3* is the RPC program version number.

System action: NFS stops.

Operator response: Check your z/OS TCP/IP setup or notify your system programmer.

System programmer response: Check your z/OS TCP/IP setup.

GFSA555I REGISTER FOR PORT *d_digits1* PROGRAM *d_digits2* - VERSION *d_digits3* - SUCCESSFUL.

Explanation: In the message text:

d_digits1 The value of *d_digits1* is the port number.

d_digits2 The value of *d_digits2* is the remote procedure call (RPC) program number.

d_digits3 The value of *d_digits3* is the RPC program version number.

System action: NFS continues processing.

GFSA556I (*procname*) UNREGISTER PROGRAM *d_digits1* VERSION *d_digits2* - FAILED.

Explanation: In the message text:

procname The name of the start procedure.

d_digits1 The value of *d_digits1* is the remote procedure call (RPC) program number.

d_digits2 The value of *d_digits2* is the RPC program version number.

System action: NFS stops.

Operator response: Contact your system programmer.

System programmer response: This message can be caused by an out-of-memory error. Increase the size of the job and step region.

GFSA557I UNREGISTER PROGRAM *d_digits1* VERSION *d_digits2* - SUCCESSFUL.

Explanation: In the message text:

d_digits1 The value of *d_digits1* is the remote procedure call (RPC) program number.

System action: NFS processing continues.

GFSA558I (*procname*) UNABLE TO CREATE IPC QUEUE.

Explanation: An operation to allocate virtual memory for an interprocess communication (IPC) queue was tried but was unsuccessful.

In the message text:

procname The name of the start procedure.

System action: NFS stops.

Operator response: Notify your system programmer.

System programmer response: Increase the size of the step region.

GFSA559I (*procname*) CANNOT CREATE UDP SERVICE.

Explanation: User datagram protocol (UDP) service transport could not be created, or you started TCP/IP before UNIX initialization was complete.

In the message text:

procname The name of the start procedure.

System action: NFS stops.

Operator response: Check your z/OS TCP/IP setup or notify your system programmer. Before starting TCP/IP, make sure that UNIX initialization is complete, and that the TCP/IP UNIX connection is established.

System programmer response: Check your z/OS TCP/IP setup. Also check the UNIX BPXPRMxx parmlib member, specifically MAXFILEPROC, MAXSOCKETS, INADDRANYPORT, and INADDRANYCOUNT. INADDRANYPORT and INADDRANYCOUNT must be specified, but their range cannot include 2049, for the server to initialize. See *z/OS MVS Initialization and Tuning Guide* for more details.

GFSA564I (*procname*) NETWORK SERVICE IS UNAVAILABLE.

Explanation: If the portmapper or TCP/IP is not running, the NFS fails to initialize and the startup ends.

In the message text:

procname The name of the start procedure.

System action: NFS startup ends.

System programmer response: Ensure that both portmapper and TCP/IP are operational.

GFSA566I (*procname*) CANNOT ACCEPT NEW TCP CLIENT CONNECTION — MAXIMUM NUMBER OF SOCKETS HAS REACHED

Explanation: When an NFS TCP client attempts to connect to the NFS server, the server cannot accept the connection because the maximum number of sockets is reached.

In the message text:

procname The name of the start procedure.

System action: The connection request fails. NFS processing continues.

Operator response: Increase the value of **maxsockets** for the AF_INET domain in the BPXPRM*xx* parmlib member so that client TCP connections can be accepted by the NFS server. See *z/OS UNIX System Services Planning* for more information.

GFSA727E NETWORK FILE SYSTEM SERVER KERBEROS DLL LOAD FAILED: *d_digits*

Explanation: Dynamic Linked Library loading of the Kerberos runtime environment has failed during the NFS server startup. This message is created when NFS is configured to require Kerberos (such that system authentication (authsys) is not allowed).

In the message text:

d_digits The value of errno after the call to the krb5_dll_load API.

System action: NFS server processing ends.

Operator response: Contact the system programmer.

System programmer response: Refer to *z/OS Integrated Security Services Network Authentication Service Programming* for more information about this failure.

GFSA728E (*procname*) SAF APPLICATION USER MAPPING FAILED WITH SAF RETURN CODE *d_digits1*, RACF RETURN CODE *d_digits2*, and RACF REASON CODE *d_digits3*.

Explanation: NFS encountered a problem in determining the RACF user ID associated with the NFS Client's Kerberos Principal.

In the message text:

d_digits1 The value of *d_digits1* is the SAF return code.

d_digits2 The value of *d_digits2* is the RACF return code.

d_digits3 The value of *d_digits3* is the RACF reason code.

System action: The NFS Client's request is failed with an RPC reply status of MSG_ACCEPTED and an accepted status of SYSTEM_ERR.

Operator response: Ensure that the NFS client principal is defined to RACF. Refer to R_usermap information in *z/OS Security Server RACF Callable*

Services for a full explanation of the SAF Return code, RACF return code and RACF reason code.

System programmer response: Refer to R_usermap information in *z/OS Security Server RACF Callable Services* for a full explanation of the SAF Return code, RACF return code and RACF reason code.

GFSA729E (*procname*) USERNAME *text1* NOT FOUND IN SAF DATABASE

Explanation: NFS encountered a problem in determining the z/OS UNIX uid/gid for this RACF user.

In the message text:

procname The name of the start procedure.

text1 The RACF User ID for which the z/OS UNIX uid/gid could not be obtained.

System action: The NFS Client's request is failed with an RPC reply status of MSG_ACCEPTED and an accepted status of SYSTEM_ERR.

Operator response: Ensure that the NFS Client Principal is correctly defined.

GFSA730I (*procname*) NETWORK FILE SYSTEM SERVER KERBEROS INITIALIZATION SUCCESSFUL

Explanation: The NFS server's Kerberos initialization was successful.

In the message text:

procname The name of the start procedure.

System action: NFS server processing continues.

GFSA731I NETWORK FILE SYSTEM SERVER KERBEROS DLL LOAD FAILED: *d_digits*

Explanation: Dynamic Linked Library loading of the Kerberos runtime environment has failed during the NFS server startup. This message is created when NFS is configured to also allow system authentication (authsys).

In the message text:

d_digits The value of the errno after the call to the krb5_dll_load API.

System action: NFS processing continues such that only requests with system authentication (sys) will be supported.

Operator response: Contact the system programmer.

System programmer response: Refer to *z/OS Integrated Security Services Network Authentication Service Programming* for more information about this failure.

GFSA733I NETWORK FILE SYSTEM SERVER KERBEROS CONTEXT CREATION FAILED: *d_digits*

Explanation: The creation of a Kerberos Context and its initialization with default values obtained from the Kerberos configuration file during the server startup has failed.

In the message text:

d_digits The value of the Kerberos error code after the call to the API krb5_init_context().

System action: NFS processing continues such that only requests with system authentication (sys) will be supported.

Operator response: Contact the system programmer.

System programmer response: Refer to *z/OS Integrated Security Services Network Authentication Service Programming* for more information about this failure.

GFSA734E NETWORK FILE SYSTEM SERVER CREDENTIAL CACHE MANAGEMENT FAILED

Explanation: There was a failure in the NFS server's Kerberos credential cache management. This could be due to one of the following reasons:

- The NFS server was unable to create a credential cache because the Kerberos API krb5_cc_generate_new() failed.
- The NFS server could not get the Kerberos principal for NFS, because the Kerberos API krb5_sname_to_principal() failed.
- The NFS server could not initialize the credential cache, because the Kerberos API krb5_cc_initialize() failed.
- The NFS server could not set the newly generated credential cache as the default cache for GSS because the GSS API gss_krb5_ccache_name() failed.

This message is created when NFS is configured to require Kerberos (i.e. system authentication (authsys) is not allowed).

System action: NFS server processing ends.

Operator response: Contact the system programmer.

System programmer response: Refer to *z/OS Integrated Security Services Network Authentication Service Programming* for more information about this failure.

GFSA735I NETWORK FILE SYSTEM SERVER CREDENTIAL CACHE MANAGEMENT FAILED

Explanation: There was a failure in the NFS server's Kerberos credential cache management. This could be due to one of the following reasons:

- The NFS server was unable to create a credential cache because the Kerberos API krb5_cc_generate_new() failed.
- The NFS server could not get the Kerberos principal for NFS, because the Kerberos API krb5_sname_to_principal() failed.
- The NFS server could not initialize the credential cache, because the Kerberos API krb5_cc_initialize() failed.
- The NFS server could not set the newly generated credential cache as the default cache for GSS because the GSS API gss_krb5_ccache_name() failed.

This message is created when NFS is configured to also allow system authentication (authsys).

System action: NFS server processing continues such that only requests with system authentication (sys) will be supported.

Operator response: Contact the system programmer.

System programmer response: Refer to *z/OS Integrated Security Services Network Authentication Service Programming* for more information about this failure.

GFSA736E NETWORK FILE SYSTEM SERVER KERBEROS TICKET ACQUISITION FAILED

Explanation: There was a failure in the NFS server's acquisition of Kerberos Ticket. This could be due to one of the following reasons:

- The NFS server was unable to read the default realm name from the Kerberos configuration file because the Kerberos API krb5_get_default_realm() failed.
- The NFS server was unable to build the TGT principal because the Kerberos API krb5_build_principal() failed.
- The NFS server could not get the Kerberos ticket using the keytab because the Kerberos API krb5_get_in_tkt_with_keytab() failed.

This message is created when NFS is configured to require Kerberos (i.e. system authentication (authsys) is not allowed).

System action: NFS server processing ends.

Operator response: Contact the system programmer.

System programmer response: Refer to *z/OS Integrated Security Services Network Authentication Service Programming* for more information about this failure.

GFSA737I NETWORK FILE SYSTEM SERVER KERBEROS TICKET ACQUISITION FAILED

Explanation: There was a failure in the NFS server's acquisition of Kerberos Ticket. This could be due to one of the following reasons:

- The NFS server was unable to read the default realm name from the Kerberos configuration file because the Kerberos API krb5_get_default_realm() failed.
- The NFS server was unable to build the TGT principal because the Kerberos API krb5_build_principal() failed.
- The NFS server could not get the Kerberos ticket using the keytab because the Kerberos API krb5_get_in_tkt_with_keytab() failed.

This message is created when NFS is configured to also allow system authentication (authsys).

System action: NFS server processing continues such that only requests with system authentication (sys) will be supported.

Operator response: Contact the system programmer.

System programmer response: Refer to *z/OS Integrated Security Services Network Authentication Service Programming* for more information about this failure.

GFSA738E NETWORK FILE SYSTEM SERVER GSS CREDENTIAL ACQUISITION FAILED

Explanation: There was a failure in the NFS server's acquisition of its GSS credentials because the GSS API gss_krb5_acquire_cred_ccache() failed. This message is created when NFS is configured to require Kerberos (i.e. system authentication (authsys) is not allowed).

System action: NFS server processing ends.

Operator response: Contact the system programmer.

System programmer response: Refer to *z/OS Integrated Security Services Network Authentication Service Programming* for more information about this failure.

GFSA739I NETWORK FILE SYSTEM SERVER GSS CREDENTIAL ACQUISITION FAILED

Explanation: There was a failure in the NFS server's acquisition of its GSS credentials because the GSS API gss_krb5_acquire_cred_ccache() failed. This message is created when NFS is configured to also allow system authentication (authsys).

System action: NFS server processing continues such that only requests with system authentication (sys) will be supported.

Operator response: Contact the system programmer.

System programmer response: Refer to *z/OS Integrated Security Services Network Authentication Service Programming* for more information about this failure.

GFSA750I (*procname*) SMF PROCESSING ACTIVE FOR USER LOGOUT.

Explanation: SMF processing is active for user logout records. This message displays in response to the **status** operand of the **modify** command.

In the message text:

procname The name of the start procedure.

System action: NFS server processing continues.

GFSA751I (*procname*) SMF PROCESSING SUSPENDED FOR USER LOGOUT.

Explanation: The system-managed facility (SMF) processing is suspended for user logout records. This message displays in response to the **status** operand of the **modify** command or after a nonzero return code from SMF. See message GFSA754I for the SMF return code.

In the message text:

procname The name of the start procedure.

System action: NFS server processing continues without generating any user or file SMF records.

Operator response: Resume SMF recording by entering a **modify** command that specifies **smf=on**.

GFSA752I (*procname*) SMF PROCESSING ACTIVE FOR FILE TIMEOUT.

Explanation: The system-managed facility (SMF) processing is active for file timeout records. This message displays in response to the **status** operand on the **modify** command.

In the message text:

procname The name of the start procedure.

System action: NFS server processing continues.

GFSA753I (*procname*) SMF PROCESSING SUSPENDED FOR FILE TIMEOUT.

Explanation: The system-managed facility (SMF) processing is suspended for file timeout records. This message displays in response to the **status** operand of the **modify** command or after a nonzero return code from SMF. See message GFSA754I for the SMF return code.

In the message text:

procname The name of the start procedure.

System action: NFS server processing continues without generating any user or file SMF records.

Operator response: Resume SMF recording by entering the **modify** command that specifies **smf=on**.

GFSA754I (*procname*) UNEXPECTED RETURN CODE *d_digits* RECEIVED FROM SMF WHILE WRITING RECORD TYPE 42 SUBTYPE [7|8].

Explanation: The NFS server received a nonzero return code, *d_digits*, while processing a file timeout (subtype 7) or user logout (subtype 8) record.

In the message text:

procname The name of the start procedure.

System action: NFS processing continues. No more SMF records of the same type and subtype are generated until the NFS address space is restarted or the **smf=on** operand of the **modify** command is entered.

Operator response: Notify your system programmer.

System programmer response: See *z/OS MVS System Management Facilities (SMF)* for information about the return code. Correct the problem, and have the operator enter a **modify** command that specifies **smf=on**.

GFSA770I z/OS UNIX REGISTRATION SUCCESSFUL.

Explanation: A connection with z/OS UNIX was established.

System action: NFS processing continues.

GFSA771I (*procname*) z/OS UNIX MOUNTS SUSPENDED.

Explanation: z/OS UNIX mount processing is suspended by the **freeze=onhfs** operand of the **modify** command. This message displays in response to either the **freeze=onhfs** operand or the **status** operand of the **modify** command. Any additional z/OS UNIX mount requests from the network are ignored. The existing mounts are unaffected.

In the message text:

procname The name of the start procedure.

System action: NFS processing continues.

GFSA772I (*procname*) z/OS UNIX MOUNTS RESUMED.

Explanation: Mount requests to the z/OS UNIXl file system have been enabled. This message displays in response to a console **modify** command.

In the message text:

procname The name of the start procedure.

System action: NFS processing continues.

GFSA776I z/OS UNIX CANNOT RESOLVE PATH NAME *text*.

Explanation: There was a failure to resolve the path name with UNIX when initializing from the mount handle data sets. The z/OS UINX file system was removed or renamed. If the user attempts to access a file object under this mount point, the NFS error response NFSERR_STALE is returned.

System action: NFS processing continues.

GFSA777I z/OS UNIX SERVICE REQUESTER DOES NOT HAVE SECURITY PRIVILEGE.

Explanation: The client user must be defined to Resource Access Control Facility (RACF) as a user of z/OS UNIX® to access z/OS UNIX file objects.

System action: The request is stopped. NFS processing continues.

System programmer response: Check the System Authorization Facility (SAF) security product user profiles.

GFSA782I (*procname*) NO ACTIVE z/OS UNIX MOUNT POINTS.

Explanation: This message is in response to the **list=mounts** operand of the **modify** command and shows that no z/OS UNIX clients are connected to NFS.

In the message text:

procname The name of the start procedure.

System action: NFS processing continues.

GFSA784I *text1* RPC z/OS UNIX ERROR VNODE_OP *text2* RC: *number1*: *text3* RSN: *number2 number3*.

Explanation: NFS encountered a problem on an interface call to z/OS UNIX. The error was encountered during processing of a *text1* remote procedure call (RPC).

In the message text:

text2 The value of *text2* data is the function called when failure occurred.

number1 The value of *number1* is the return code.

text3 The value of *text3* text is the English description of the return code.

number2 number3 The combination of *number2 number3* represents the reason code as returned by z/OS UNIX.

System action: NFS processing continues.

System programmer response: See z/OS UNIX documentation for a full explanation of z/OS UNIX reason codes (for example, *z/OS UNIX System Services Messages and Codes*).

GFSA786I MULTI-COMPONENT LOOKUP REQUEST FOR PATHNAME *text* CANNOT BE RESOLVED.

Explanation: The multicomponent **lookup** request for the path name failed. Possible reasons follow:

- Symbolic links cannot be embedded in a multicomponent path name because it is not supported at this time.
- Either the specified path name specified is not supported, or access is not allowed.
- A public path is not set up on this server.

System action: The request fails. NFS processing continues.

User response: Construct a different **lookup** request with a valid path name.

GFSA787I z/OS UNIX PATHNAME SPECIFIED, BUT NOHFS SPECIFIED IN THE ATTRIBUTE DATA SET.

Explanation: A z/OS UNIX path name was specified but **nohfs** was also specified in the attributes data set, which disables z/OS UNIX processing.

System action: Network File System Server request fails.

User response: Either change the attributes data set to activate z/OS UNIX processing and restart the NFS server, or change the path name.

GFSA791I (*procname*) Owner (*serverid clientid userid processid*) offset=*hex1* len=*hex2* lockacc=*mode* status=*status*

Explanation: This message displays a byte range lock owner for the file requested in the LISTLOCKS operator command. Message GFSA791I will be returned for each lock. The lock owner is identified by *serverid*, *clientid*, *userid*, and *processid*. The lock is for offset *hex1* and length *hex2*. The lock access *mode* (share/exclusive) and the current lock *status* (waiting/granted) are also displayed.

System action: The Network File System server continues.

Operator response: None.

System programmer response: None.

Source: z/OS NFS Server

GFSA792I **(***procname***) Owner (***serverid clientid userid processid***) access=***access* **deny=***deny*

Explanation: This message displays a lock share holder for the file requested in the LISTLOCKS operator command. Message GFSA792I will be returned for each share holder. The share owner is identified by *serverid*, *clientid*, *userid*, and *processid*. The lock access mode *access* and deny mode *deny* are displayed.

System action: The Network File System server continues.

Operator response: None.

System programmer response: None.

Source: z/OS NFS Server

GFSA793I **(***procname***)** *text* **does not have Locks.**

Explanation: This message is returned from the LISTLOCKS operator command for the specified file *text*, if the file does not have any locks or shares.

System action: The Network File System server continues.

Operator response: None.

System programmer response: None.

Source: z/OS NFS Server

GFSA794I **(***procname***) LISTLOCK** *text* **command was completed.**

Explanation: This message indicates that processing of the LISTLOCKS operator command for the specified file *text* has been completed. Any existing locks will be listed in message GFSA794I.

System action: The Network File System server continues.

Operator response: None.

System programmer response: None.

Source: z/OS NFS Server

GFSA801I **MOUNT FAILED:** *text*

Explanation: The value of *text* can be any of these messages.

- `FILE MAPPING ENABLED BUT NO SIDE FILE SPECIFIED`
- `STORAGE LIMIT REACHED LOADING MAPPING SIDE FILE`
- `MAPPING SIDE FILE NOT FOUND`
- `ERROR OPENING/READING MAPPING SIDE FILE`
- `MAPPING SIDE FILE HAS INVALID SYNTAX OR FORMAT`
- `SIDE FILE SPECIFIED BUT MAPPING IS DIS-ALLOWED BY INSTALLATION`

User response: Take one of these actions depending on the value of *text*:

- Specify a side file if **fileextmap** is ON.
- Fix the problem with mapping the side file.
- Ask the system administrator to change the **sfmax** value, and then reissue the **mount** command.

GFSA802E **REMOUNT FAILED — PHYSICAL FILE SYSTEM CHANGED. PATH: pathname PREV: datasetname CURRENT: datasetname**

Explanation: During a restart of the NFS server, the rebuild of the mount point recorded in the mount handle database failed. The failure occurred because the physical file system for the mount point was changed by a TSO UNMOUNT command because the mount point was originally mounted.

System action: NFS server processing continues.

GFSA811I **CANNOT FLUSH PARTIAL RECORDS FOR DATA SET** *text1***(***text2***): FB** *h_digits***.**

Explanation: There was not enough memory to allocate the storage required to flush partial records to the data set at data set close time. The partial record was discarded.

System action: The data set is closed. NFS processing continues.

System programmer response: Increase the size of the step region. The data set might be incomplete.

GFSA812I **FLUSH FAILED: RC** *d_digits1* **OFFSET** *d_digits2* **WAS DROPPED FOR DATA SET** *text1***(***text2***).**

Explanation: When NFS attempted to flush cached data at data set close time, the error *d_digits1* was detected. The data at offset *d_digits2* was discarded. This error message follows more specific error messages.

System action: The data set is closed. NFS processing continues.

System programmer response: See the message preceding this message to determine the correct action. The data set might be incomplete. See Table 61 on page 299 for a description of the return code *d_digits1*.

GFSA813I **REMOVE FAILED: RC** *h_digits* **DSN** *text1***(***text2***).**

Explanation: The error *h_digits* was detected during an attempt to remove member *text2* from PDS *text1*.

System action: NFS processing continues.

System programmer response: The value of *h_digits* is the return code from the z/OS DFSMS STOW macro.

See *z/OS DFSMS Macro Instructions for Data Sets* for the explanation of STOW return codes.

GFSA814I RENAME FAILED: RC *h_digits* **DSN** *text1* **OLDMEM** *text2* **NEWMEM** *text3***.**

Explanation: The error *h_digits* was detected during an attempt to rename member *text2* to *text3* in partitioned data set (PDS) *text1*.

System action: NFS processing continues.

System programmer response: The value of *h_digits* is the return code from the z/OS DFSMS STOW macro. See *z/OS DFSMS Macro Instructions for Data Sets* for the explanation of STOW return codes.

GFSA815I RENAME FAILED: RC *h_digits1* **OLDDSN** *text1* **VOL** *text2* **UNIT** *h_digits2* **NEWDSN** *text3***.**

Explanation: The error *h_digits1* was detected during an attempt to rename the non-VSAM data set *text1* to *text3*. The value of *text2* is disk volume serial number. The value of *h_digits2* is the disk device type code.

System action: NFS processing continues.

System programmer response: The value of *h_digits1* is a composite of one or more error return codes that the system encountered when trying to rename the data set. You can decode the hexadecimal digits using the following list.

Renaming a data set requires three steps:

1. Uncatalog the old data set name,
2. Rename the data set in the disk VTOC, and
3. Catalog the new data set name.

Should an error occur in either of the last two steps, the prior step or steps are undone to preserve the old data set name.

Find the step that failed by matching the value in byte zero of the return code with one of the values under the heading **Byte 0** in the following list. Byte three contains the return code from the first failing DFP service (uncatalog/catalog or DADSM rename). If more errors occur during an attempt to recatalog or rename the data set back to the old name, the return codes are placed in bytes one and two respectively.

Byte 0	Meaning/other bytes
00	Error uncataloging old data set name. Byte 3: Uncatalog return code.
01	Error renaming the data set. Byte 1: Recatalog return code for old data set name. Byte 3: DADSM rename return code.
02	Error cataloging new data set name. Byte 1: Recatalog return code for old data set name. Byte 2: DADSM rename return code for old data set name. Byte 3: Catalog return code for new data set name.

Catalog and DADSM rename return codes are documented in *z/OS DFSMSdfp Advanced Services*.

Message IEC614I is written for DADSM rename errors and contains more diagnostic codes. These codes are documented in *z/OS DFSMSdfp Diagnosis*.

GFSA816I HOST NAME OF IP ADDRESS (*d_digits***) WAS NOT FOUND BY TCP/IP.**

Explanation: The client host name of IP address *d_digits* is not defined in either the TCP/IP domain name server or the TCP/IP site table. See *TCP/IP: Performance Tuning Guide*.

System action: NFS processing continues. The dotted IP address is used as the host name.

System programmer response: Insert this client host's entry into either the TCP/IP domain name server or the TCP/IP site table.

GFSA817I *text1* **REQUEST NOT VALID ON ALIAS NAME** *text2***.**

Explanation: The value of *text1* can be remove or rename. The value of *text2* is an MVS access method services alias name of a file that also has a true name. Remove (**rm** or **rmdir**) and rename (**mv**) requests cannot be run using an alias name. The true file name is required.

System action: The request is stopped. An I/O error indication is returned to the client. NFS processing continues.

System programmer response: Inform the client user of this error.

User response: Provide the true name of the file in the request.

GFSA818I (*procname***) EXPORTS: NO VALID HOST NAMES IN** *text* **LIST.**

Explanation: None of the client host names in the read/write or access list is defined to the network (see *text*).

In the message text:

procname The name of the start procedure.

System action: The NFS does not export the associated directory if the result is a null access list. If the result is a null read/write list, the directory is exported with read-only access.

System programmer response: Correct the host names in the exports data set or have the host names defined to the network.

GFSA819I DATA SET *text1* CREATION USING DATA CLASS = *text2*.

Explanation: Data set *text1* is being allocated using the attributes in data class *text2*.

System action: NFS processing continues.

GFSA820I CATALOG ERROR OCCURRED WHILE [RETRIEVING|UPDATING] CATALOG INFORMATION FOR *text*. RETURN CODE IS *d_digits1*, REASON CODE IS *cc-d_digits2*.

Explanation: The catalog management module IGG0CL*cc* returned the return code *d_digits1* and reason code *d_digits2* as the result of a catalog error or an exception condition. The value of *text* is the name of the data set against which the retrieve or update operation was performed.

System action: NFS processing continues.

System programmer response: See message IDC3009I in *z/OS MVS System Messages, Vol 6 (GOS-IEA)* for specific return code and reason code information.

GFSA821I ERROR OCCURRED WHILE UPDATING THE FORMAT 1 DSCB FOR *text1* ON *text2*. FUNCTION CODE IS *d_digits1*, RETURN CODE IS *d_digits2*, REASON CODE IS *d_digits3*.

Explanation: In the message text:

text1 The value of *text1* is the name of the data set.

text2 The value of *text2* is the serial number of the volume on which the data set resides.

The function code can be one of the following codes:

2 Deserializing the unit control block (UCB)

4 Deserializing the direct access storage device (DASD) volume

12 Searching for the UCB

16 Serializing the DASD volume

20 Reading the data set control block (DSCB)

24 Writing the DSCB

System action: NFS processing continues.

System programmer response: See the following manuals for specific return code and reason code information:

2 *z/OS MVS Programming: Authorized Assembler Services Reference SET-WTO*, macro UCBPIN

4 *z/OS MVS Programming: Authorized Assembler Services Reference ALE-DYN*, macro DEQ

12 *z/OS MVS Programming: Authorized Assembler Services Reference SET-WTO*, macro UCBLOOK

16 *z/OS MVS Programming: Authorized Assembler Services Reference LLA-SDU*, macro RESERVE

20 *z/OS DFSMSdfp Advanced Services*, return codes from CVAFDIR

24 *z/OS DFSMSdfp Advanced Services*, return codes from CVAFDIR

GFSA822I (*procname*) UNABLE TO PERFORM FILE MAPPING BECAUSE NO SIDE FILE SPECIFIED OR LOADED

Explanation: File mapping cannot be performed because a side file was not specified either as a default or at the mount point.

In the message text:

procname The name of the start procedure.

System action: The operation fails.

System programmer response: Specify a side file in the attributes data set.

User response: Specify a side file in the **mount** command.

GFSA823I (*procname*) PUBLIC PATH CANNOT BE ESTABLISHED.

Explanation: The public path name(s) specified in the **public** keyword cannot be established during server startup. This could be because the path is not exported or because it does not exist.

In the message text:

procname The name of the start procedure.

System action: NFS server startup ends.

Operator response: Make sure that the public path name exists and is exported if export list checking is enabled. Correct the problem and restart the server.

GFSA827I REQUEST *h_digits* CREDENTIALS ALLOCATION FAILED.

Explanation: An operation to allocate virtual memory for a credentials block was attempted but was unsuccessful.

In the message text:

h_digits The value of *h_digits* is the request block address.

System action: The request is stopped. NFS processing continues.

System programmer response: Increase the size of the step region.

GFSA829I REQUEST *h_digits* INVALID CREDENTIALS FLAVOR *d_digits*.

Explanation: An incorrect credentials type was received from the client. This is probably a client software error. NFS supports UNIX and nonauthentication styles.

In the message text:

h_digits The value of *h_digits* is the request block address.

System action: The request is stopped. NFS processing continues.

System programmer response: Inform the client user that the credentials used are not valid.

GFSA832I REQUEST *h_digits* INVALID MEMBERNAME FOR *text*.

Explanation: In the message text:

text The value of *text* is the member name of a partitioned data set (PDS) that was specified as a file name by the NFS client user. The file name specified was incorrect or was not found in the PDS.

h_digits The value of *h_digits* is the request block address.

System action: The request is stopped. NFS processing continues.

System programmer response: Inform the client user of this error.

User response: Correct the error and resubmit the request.

GFSA833I REQUEST *h_digits* PARSE FAILED FOR *text*.

Explanation: In the message text:

text The value of *text* is the member name of a partitioned data set (PDS) or a data set name that was specified as a file name by the NFS client user. The file name specified was incorrect, was not found in the PDS, or was an incorrect or nonexistent data set.

h_digits The value of *h_digits* is the request block address.

System action: The request is stopped. NFS processing continues.

System programmer response: Inform the client user of this error.

User response: Correct the error and resubmit the request.

GFSA840I DYNAMIC ALLOCATION FAILED WITH RETURN CODE *h_digits* FOR DATA SET *text(dsname)*.

Explanation: A dynamic file allocation error occurred.

In the message text:

h_digits The value of *h_digits* is the dynamic allocation return code.

dsname The value of *dsname* is the data set name.

text The value of *text* is the member name, if any.

System action: The request is stopped. NFS processing continues.

System programmer response: This message is preceded by either message GFSA853I or GFSA854I. See the programmer response for the message that precedes this message to determine the appropriate action.

GFSA842I *dsname1* UNSUPPORTED DSORG *dsname2*.

Explanation: NFS does not support data set organization of data set *dsname1*. The value of *dsname2* can be `ISAM` or `UNKNOWN`.

System action: The request is stopped. NFS processing continues.

System programmer response: Inform the client user of this error.

GFSA843I CREATE FAILED FOR *dsname*.

Explanation: An error occurred while the NFS was trying to create the data set *dsname*. This message follows other messages that describe the error in greater detail.

System action: The request is stopped. NFS processing continues.

System programmer response: See the message(s) preceding this message to determine the appropriate response.

GFSA847I IDCAMS ERROR: *text*.

Explanation: The value of *text* is an access method services error message.

System action: The request is stopped. NFS processing continues.

System programmer response: See *z/OS MVS System Messages* for more information about the access method services error message, or use LookAt. For a description of LookAt, see "How to look up message explanations with LookAt" on page xv.

GFSA848I PDS *text* IS NOT EMPTY.

Explanation: The NFS client user issued a **rmdir** (remove directory) AIX or UNIX command to remove a partitioned data set (PDS) that was not empty. The NFS version 2 protocol specification requires the directory (PDS) to be empty before it is removed. This is a NFS client user error.

In the message text:

text The value of *text* is the name of the PDS.

System action: The request is stopped. An error is returned to the client. NFS processing continues.

System programmer response: Inform the NFS client user of this error.

User response: Remove all files in the directory, and then resubmit the **rmdir** request.

GFSA849I NEW AND OLD FILES ARE NOT MEMBERS OF THE SAME PDS.

Explanation: Rename is not allowed for a member of a partitioned data set (PDS) when the target name is not in the same PDS. This is a NFS client user error.

System action: The request is stopped. An error is returned to the client. NFS processing continues.

System programmer response: Inform the NFS client user of this error.

User response: Check the file name used. Correct it and try again.

GFSA853I DYNAMIC ALLOCATION: INPUT VALIDATION ROUTINE REJECTED ALLOCATION.

Explanation: Dynamic allocation failed when issued by the installation input validation routine. Installation configuration errors probably occurred.

System action: The request is stopped. NFS processing continues.

System programmer response: See the *z/OS MVS Installation Exits* for more information about the Input Validation routine.

GFSA854I DYNAMIC ALLOCATION: RC=*d_digits* ERROR=*h_digits1* INFO=*h_digits2* : *text*.

Explanation: Dynamic allocation failed with return code *d_digits*, error reason code *h_digits1*, and information reason code *h_digits2*.

In the message text:

text The value of *text* is a description of the interrupted dynamic allocation request.

System action: The request is stopped. NFS processing continues.

System programmer response: See *z/OS MVS Programming: Authorized Assembler Services Guide* for more information about these codes.

GFSA858I OPEN FAILED RC *d_digits* FOR DATA SET *dsname1*(*dsname2*).

Explanation: An error occurred during an attempt to open the data set.

In the message text:

dsname1 The value of *dsname1* is the data set name.

dsname2 The value of *dsname2* is the member name, if any.

d_digits The value of *d_digits* is the return code.

System action: The request is stopped. NFS processing continues.

System programmer response: If this error was not caused by an out-of-memory condition, contact your programming support personnel. See Table 61 on page 299 for a description of the return code *d_digits*.

GFSA859I READ FAILED RC *d_digits* FOR DATA SET *dsname1*(*dsname2*).

Explanation: An error occurred during an attempt to read the data set.

In the message text:

dsname1 The value of *dsname1* is the data set name.

dsname2 The value of *dsname2* is the member name, if any.

d_digits The value of *d_digits* is the return code.

System action: The request is stopped. NFS processing continues.

System programmer response: If this error was not caused by an out-of-memory condition, contact your programming support personnel. See Table 61 on page 299 for a description of the return code *d_digits*.

GFSA860I WRITE FAILED RC *d_digits* FOR DATA SET *dsname1*(*dsname2*).

Explanation: An error occurred during an attempt to write the data set.

In the message text:

dsname1 The value of *dsname1* is the data set name.

dsname2 The value of *dsname2* is the member name, if any.

d_digits The value of *d_digits* is the return code.

System action: The request is stopped. NFS processing continues.

System programmer response: If this error was not caused by an out-of-memory condition, contact your programming support personnel. See Table 61 on page 299 for a description of the return code *d_digits*.

GFSA862I CATALOG (*text*) COULD NOT BE LOCATED.

Explanation: The user catalog named *text* which contains the entry for an index, could not be located. The catalog does not exist or is not mounted. If it does not exist, the entry in the master catalog might be incorrect.

System action: NFS processing continues.

System programmer response: Investigate why the catalog could not be found and take corrective action.

GFSA863I READDIR ON ROOT IS NOT ALLOWED.

Explanation: The user attempted to list the contents of the master catalog.

System action: The request is stopped. NFS processing continues.

GFSA864I (*procname*) CANNOT OPEN THE EXPORTS DATA SET.

Explanation: The server was unable to open the exports data set defined in the job control language (JCL) for DDNAME EXPORTS. The DD statement might be missing or the data set name might be incorrect.

In the message text:

procname The name of the start procedure.

System action: NFS stops.

Operator response: Notify your system programmer.

System programmer response: Correct the JCL for DDNAME EXPORTS.

GFSA865I (*procname*) EXPORTS: UNEXPECTED OPTION (*text*)-- SHUTDOWN SCHEDULED.

Explanation: The option information provided in the *text* data is incorrect. This error could occur as a result of unexpected blanks, incorrect syntax, or mutually exclusive options (for example, both **ro** and **rw**).

In the message text:

procname The name of the start procedure.

System action: Checking of the exports data set continues, but the shutdown of NFS occurs at its completion.

System programmer response: Correct the exports data set and restart NFS.

GFSA866I (*procname*) EXPORTS: DIRECTORY *dsname* WAS NOT EXPORTED.

Explanation: An error was encountered that was severe enough to prevent the data set or index named *dsname* from being exported. This message follows a more specific error message.

In the message text:

procname The name of the start procedure.

System action: NFS processing continues.

System programmer response: Correct the exports data set.

GFSA867I (*procname*) EXPORTS: *dsname1* CANNOT BE EXPORTED BECAUSE *dsname2* ALREADY IS.

Explanation: The data set or index named in *dsname1* is a parent directory or a subdirectory of the data set or index named in *dsname2* which is already exported.

In the message text:

procname The name of the start procedure.

System action: NFS processing continues.

System programmer response: Correct the exports data set.

GFSA868I (*procname*) EXPORTS: HOST (*text*) WITH UNKNOWN ADDRESS TYPE (*d_digits*) HAS BEEN IGNORED.

Explanation: The host named *text* has an IP address (*d_digits*) that is not valid.

In the message text:

procname The name of the start procedure.

System action: NFS processing continues.

System programmer response: Correct the exports data set.

GFSA869I (*procname*) EXPORTS: UNKNOWN HOST (*dsname*) HAS BEEN IGNORED.

Explanation: The client host named *dsname* is not defined to the network. If other defined client hosts are specified for the option, NFS ignores this undefined client host as if had not been specified. (However, see message GFSA818I for special handling of null lists.)

In the message text:

procname The name of the start procedure.

System action: NFS processing continues.

System programmer response: Correct the host name in the exports data set or have this host name defined to the network.

GFSA871I REQUEST *h_digits* **HAS MISMATCHED UID: CRED =** *d_digits1* **ARGS =** *d_digits2***.**

Explanation: The value of *h_digits* is a block created for the logon or logout request. The value of *d_digits1* represents the credential user ID number. The value of *d_digits2* represents the client user ID number. The user ID numbers do not match, and this is considered a security failure.

System action: The client logon or logout request is stopped. NFS processing continues.

System programmer response: This is a NFS client application problem. If the NFS client application is offered by IBM, inform the IBM Support Center. If the NFS client application is offered by the programming support personnel, contact your programming support personnel.

GFSA876I I/O ERROR ON DSN = *text1***(***text2***) SENSE** *h_digits1* **IOBCSW** *h_digits2 h_digits3***. ACCESS METHOD RC =** *h_digits4* **ACCESS METHOD RSN =** *h_digits5*

Explanation: The physical I/O layer tried to check some previous operation in the data set and the check failed.

In the message text:

text1 The value of *text1* is the data set name.

text2 The value of *text2* is the member name, if any.

h_digits1 The value of *h_digits1* is the sense bytes 0 and 1 from the device.

h_digits2 The value of *h_digits2* is the first 3 bytes of the channel status word from the device.

h_digits3 The value of *h_digits3* is the last 4 bytes of the channel status word from the device.

h_digits4 The value of *h_digits4* is the access method return code.

h_digits5 The value of *h_digits5* is the access method reason code.

System action: The request is stopped. NFS processing continues.

System programmer response: See the appropriate device documentation for more information on the sense bytes and channel status word.

GFSA877I R0=*h_digits1* **R1=***h_digits2***:** *text* **ACCESS METHOD RC =** *h_digits3* **ACCESS METHOD RSN =** *h_digits4*

Explanation: A SYNAD error was detected during a physical I/O operation.

In the message text:

h_digits1 The value of *h_digits1* is the contents of register 0.

h_digits2 The value of *h_digits2* is the contents of register 1.

text The value of *text* is the message returned from the SYNAD analysis function macro.

h_digits3 The value of *h_digits3* is the access method return code.

h_digits4 The value of *h_digits4* is the access method reason code.

System action: The request is stopped. If the error detected is a B37, D37, or E37 abend, NFS restores the file size to the last known file size before the SYNAD error. NFS processing continues.

System programmer response: A data management message should have displayed on the console. See *z/OS MVS System Messages, Vol 1 (ABA-AOM)* through *z/OS MVS System Messages, Vol 9 (IGF-IWM)* for a description of the return code to determine the corrective action, or use LookAt. For a description of LookAt, see "How to look up message explanations with LookAt" on page xv.

GFSA878I (*procname***) CANNOT OPEN THE CHKLIST DATA SET.**

Explanation: NFS cannot open the checklist data set, specified by the CHKLIST DD statement in the MVSNFS startup procedure.

In the message text:

procname The name of the start procedure.

System action: NFS ends.

Operator response: Contact the system programmer.

System programmer response: Check for validity of the specified data set. Correct the JCL for the CHKLIST DD statement in the MVSNFS startup procedure.

GFSA879I (*procname*) PROBLEMS ENCOUNTERED PARSING THE CHKLIST DATA SET.

Explanation: NFS was not able to parse the checklist data set specified by the CHKLIST DD statement successfully.

In the message text:

procname The name of the start procedure.

System action: NFS ends.

Operator response: Contact the system programmer.

System programmer response: Correct the contents of the checklist data set.

GFSA881I FUB: *h_digits1* UNABLE TO ACCESS FILE *text1*(*text2*) OWNED BY FUB *h_digits2*.

Explanation: A user tried to access a data set already locked internally by NFS for writing by another user. The data set has not been released yet.

In the message text:

h_digits1 The value of *h_digits1* is the address of the file usage block.

text1 The value of *text1* is the data set name.

text2 The value of *text2* is the member name, if any.

h_digits2 The value of *h_digits2* is the address of the file usage block that currently has the file allocated.

System action: The request is stopped. A `Not Owner` error message is returned to the user. NFS processing continues.

System programmer response: If queried by the user, the `Not Owner` error message, as it relates to NFS, is described in *z/OS Network File System Guide and Reference*.

GFSA886I UNABLE TO WRITE RECNO *d_digits* TO *text1*(*text2*) DUE TO PREVIOUS ERROR.

Explanation: An error was previously detected which prevents writing to a data set.

System action: The request is stopped. NFS processing continues.

System programmer response: If the previous error cannot be determined, contact your programming support personnel.

GFSA895I REQUEST *h_digits* - FILE *text* NOT ALLOCATED.

Explanation: NFS did not have the data set *text* open during a request to close the file. The file name might have been specified incorrectly, or the timeout might have already occurred for this data set, causing the server to close the data set.

In the message text:

h_digits The request block address.

text The data set name.

System action: The system ends the request. Return code X'131' is passed back to the user. The system processing continues for NFS.

GFSA896I REQUEST *h_digits1* - FILE BLOCK *h_digits2* ASSOCIATED WITH FILE *text* NOT IN USE BY CREDENTIALS *h_digits3*.

Explanation: The request to close file *text* was received, but the file was not opened by the client. A file can be closed only by the same client that opened the file.

In the message text:

h_digits1 The value of *h_digits1* is the request block address.

h_digits2 The value of *h_digits2* is the file block address.

h_digits3 The value of *h_digits3* is the credentials block address.

System action: The request is stopped. Return code 132 is passed back to the client. NFS processing continues.

GFSA897I RECORD *d_digits1* SIZE *d_digits2* IS *text1*: *text2* = *d_digits3* DSN = *text3*(*text4*).

Explanation: The record received from the client can be too long or too short.

In the message text:

d_digits1 The value of *d_digits1* is the record number.

d_digits2 The value of *d_digits2* is the record size.

text1 The value of *text1* can be `TOO LONG` or `TOO SHORT`.

text2 The value of *h_digits1* can be `MAXIMUM` or `MINIMUM`.

d_digits3 The value of *h_digits1* shows the maximum or minimum length allowed.

text3 The value of *text3* is the data set name

text4 The value of *text4* is the member name, if any.

System action: The request is stopped. NFS processing continues.

GFSA898I EOL SEQUENCE MISMATCH FOR DATA SET *text1*(*text2*).

Explanation: The end-of-line terminator was not found in the same place as the previous end-of-line terminator for an offset that is being rewritten by the client. This is a NFS client error.

System action: The request is stopped. NFS processing continues.

User response: Make sure that the record you are writing is the same size as the record you are replacing.

GFSA899I BLANKSTRIP MODE: TRAILING BLANK(S) IN RECORD *d_digits* IS NOT ALLOWED DSN = *text1*(*text2*).

Explanation: Writing data in text mode with blank stripping enabled and blanks at the end of the line to a data set with fixed-length records is not allowed.

System action: The request is stopped. NFS processing continues.

GFSA900I (*procname*) MOUNT PROCESSING ACTIVE.

Explanation: Mounts can be issued from the network. This message displays in response to the **status** operand of the **modify** command.

In the message text:

procname The name of the start procedure.

System action: NFS processing continues.

GFSA901I (*procname*) MOUNT PROCESSING SUSPENDED.

Explanation: Mount processing was suspended by the **freeze=on** operand of the **modify** command. This message displays in response to either the **freeze=on** operand or the **status** operand of the **modify** command. Any additional mount requests from the network are ignored. Existing mounts are unaffected.

In the message text:

procname The name of the start procedure.

System action: NFS processing continues.

GFSA902I (*procname*) MOUNT PROCESSING RESUMED.

Explanation: Mount processing has been resumed by the **freeze=off** operand of the **modify** command. Any additional mount requests from the network are honored. Existing mounts are unaffected.

In the message text:

procname The name of the start procedure.

System action: NFS processing continues.

GFSA903I (*procname*) TASK *h_digits1* TCB *h_digits2* PROGRAM = *text1* = *text2*.

Explanation: This is in response to the **status** operand of the **modify** command.

In the message text:

procname The name of the start procedure.

h_digits1 The value of *h_digits1* is the task queue address.

h_digits2 The value of *h_digits3* is the task control block address.

text1 The value of *text1* is the program name.

text2 The value of *text2* is the specific task name.

System action: NFS processing continues.

GFSA904I (*procname*) z/OS UNIX MOUNT PROCESSING ACTIVE.

Explanation: Mounts can be issued from the network. This message displays in response to the **status** operand of the **modify** command.

In the message text:

procname The name of the start procedure.

System action: NFS processing continues.

GFSA905I (*procname*) Current CTRACE buffer is flushed and switch made to the next buffer.

Explanation: In response to the FLUSHCTR operator command the Network File System Component Trace function flushed the remaining trace buffers to the component trace external writer and switched to the next available buffer.

System action: The Network File System server continues.

Operator response: None.

System programmer response: None.

Source: z/OS NFS Server

GFSA908I (*procname*) z/OS UNIX PROCESSING DISABLED.

Explanation: z/OS UNIX file system processing is suspended. This message displays in response to the **status** operand of the **modify** command.

In the message text:

procname The name of the start procedure.

System action: NFS processing continues.

GFSA909E (*procname*) UNMOUNT COMMAND FAILED: MOUNT POINT STILL IN USE

Explanation: This message is in response to the **unmount** operand of the **modify** command. The unmount processing fails because the file system is still in use. The user can retry the **unmount** command later after there is no reference to the file system.

In the message text:

procname The name of the start procedure.

System action: The system ends the request. The system processing continues for NFS.

GFSA910I (*procname*) *text* ACTIVE = *d_digits*.

Explanation: This message is in response to the **list=mounts** operand of the **modify** command and shows the mounted file space (MVS path) and the number of active clients (*d_digits*) mounting those files by that path. The number might not be accurate if a client crashed without unmounting the files.

In the message text:

procname The name of the start procedure.

System action: NFS processing continues.

GFSA911I (*procname*) *text1*(*text2*) ACTIVE = *d_digits*.

Explanation: This status message is in response to the **list=mounts** operand of the **modify** command and shows the mounted file space (MVS path) and the number of active clients (*d_digits*) mounting those files by that path. The number might not be accurate if a client crashed without unmounting the files.

In the message text:

procname The name of the start procedure.

System action: NFS processing continues.

GFSA912I (*procname*) *dsname*.

Explanation: The data set name *dsname* appears in response to a **list=dsnames** operand of the **modify** command and shows a currently active data set.

In the message text:

procname The name of the start procedure.

System action: NFS processing continues.

GFSA913I (*procname*) *text1*(*text2*).

Explanation: The PDS member appears in response to the **list=dsnames** operand of the **modify** command and shows a currently active data set.

In the message text:

procname The name of the start procedure.

System action: NFS processing continues.

GFSA914I (*procname*) *dsname* DEALLOCATED.

Explanation: The data set name *dsname* appears in response to a **release=dsname(member)** operand of the **modify** command after successful deallocation.

In the message text:

procname The name of the start procedure.

System action: NFS processing continues.

GFSA915I (*procname*) *dsname* NOT ALLOCATED.

Explanation: The data set name *dsname* appears in response to the **release=dsname(member)** operand of the **modify** command if the data set or member specified to be released was not found.

In the message text:

procname The name of the start procedure.

System action: NFS processing continues.

GFSA916I (*procname*) *dsname* UNMOUNTED.

Explanation: The data set name *dsname* appears in response to an **unmount** (data set or member) operand of the **modify** command after a successful unmount.

In the message text:

procname The name of the start procedure.

System action: NFS processing continues.

GFSA917I (*procname*) *dsname* NOT MOUNTED.

Explanation: The data set name *dsname* appears in response to the **unmount** (data set or member) operand of the **modify** command if the data set or member specified to be unmounted was not found in the current mount list.

In the message text:

procname The name of the start procedure.

System action: NFS processing continues.

GFSA918I (*procname*) *dsname* IS NOT A VALID DATA SET NAME.

Explanation: The data set name *dsname* specified in either the **release** operand or the **unmount** operand of the **modify** command is not a valid MVS data set name.

In the message text:

procname The name of the start procedure.

System action: NFS processing continues.

Operator response: Specify the command again with a valid data set name.

GFSA919I (*procname*) *text* IS NOT A VALID MEMBER NAME.

Explanation: The member name *text* specified in either the **release** operand or the **unmount** operand of the **modify** command is not a valid MVS member name.

In the message text:

procname The name of the start procedure.

System action: NFS processing continues.

Operator response: Specify the command again with a valid member name.

GFSA920I (*procname*) NO ACTIVE MOUNT POINTS.

Explanation: This message is in response to the **list=mounts** operand of the **modify** command and shows that there are no clients connected to the NFS.

In the message text:

procname The name of the start procedure.

System action: NFS processing continues.

GFSA921I (*procname*) NO ACTIVE DATA SETS.

Explanation: This message is in response to the **list=dsnames** operand of the **modify** command and shows that there are no clients actively accessing data sets.

In the message text:

procname The name of the start procedure.

System action: NFS processing continues.

GFSA922I (*procname*) VERIFY: (*dsname*) IS NOT A VSAM DATA SET.

Explanation: This message is in response to the **verify** operand of the **modify** command and shows that the data set name *dsname* specified is not a virtual storage access method (VSAM) data set.

In the message text:

procname The name of the start procedure.

dsname The value of *dsname* specifies the data set name.

System action: The system ends the request. The system processing continues for the NFS.

GFSA923I (*procname*) VERIFY SUCCESSFUL FOR (*dsname*).

Explanation: This message is in response to the **verify** operand of the **modify** command. It shows that the verify was successful for the VSAM data set.

In the message text:

procname The name of the start procedure.

dsname The value of *dsname* specifies the data set name.

System action: The system processing continues for NFS.

GFSA924I (*procname*) VERIFY FAILED WITH RC = *d_digits* FOR (*dsname*).

Explanation: This message is in response to the **verify** operand of the **modify** command. It shows that the **verify** operand failed with a return code for the data set. Message GFSA847I follows this message in the log data set.

In the message text:

procname
The name of the start procedure.

d_digits The value of *d_digits* is the return code.

dsname The value of *dsname* is the data set name.

System action: The system ends the request. The system processing continues for NFS.

GFSA925I (*procname*) ERROR WAS DETECTED IN THE EXPORTS FILE. EXPORT LIST NOT REBUILT.

Explanation: This is the response from the **exportfs** operand of the **modify** command, indicating that one or more errors were detected in the exports data set (for example, the exports data set cannot be opened).

In the message text:

procname The name of the start procedure.

System action: NFS processing continues. The existing exports list is not changed.

Operator response: Notify your system programmer.

System programmer response: Review previous console error messages for detailed information about the specific error in the exports data set.

GFSA926I (*procname*) EXPORT LIST HAS BEEN REBUILT SUCCESSFULLY.

Explanation: This is the reply from the **exportfs** operand of the **modify** command, indicating that the command completed normally.

In the message text:

procname The name of the start procedure.

System action: NFS processing continues with the new exports data set in effect.

GFSA927I (*procname*) MODIFY EXPORTS COMMAND IGNORED - THE EXPORTS FILE IS NOT BEING USED FOR SECURITY CHECKING.

Explanation: This is the reply from the **exportfs** operand of the **modify** command, indicating that the command was ignored because the site attribute for security requested that no exports file checking be done.

In the message text:

procname The name of the start procedure.

System action: NFS processing continues with the security options unchanged.

GFSA928I MAPFILE: MODIFY MAPFILE COMMAND WAS COMPLETED SUCCESSFULLY.

Explanation: This is the reply from the **mapfile** operand of the **modify** command, indicating the command completed normally.

System action: NFS processing continues with the new side files data set in storage.

GFSA929I MODIFY MAPFILE COMMAND IGNORED - MAPPED SIDEFILES ARE NOT BEING USED.

Explanation: This is the reply from the **mapfile** operand of the **modify** command, indicating that the command was ignored because mapped processing attribute was not specified in the site attribute file (NFSATTR) or any **mount** command.

System action: NFS processing continues with the side files data unchanged.

GFSA930I (*procname*) LOG DATA SET IS SWITCHED FROM *text1* TO *text2*.

Explanation: A "no space" or an I/O error condition is detected when writing to the log data set. NFS logging is now switched to the other log data set.

In the message text:

procname
The name of the start procedure.

text1 The value of *text1* is the data definition (DD) associated with the switched-from log data set.

text2 The value of *text2* is the DD associated with the switched-to log data set.

System action: NFS processing continues.

Operator response: If requested by the installation, back up the switched-from log data set at this point. The switched-from data set is reused when the switched-to log data set is also filled.

System programmer response: Consider allocating larger NFS server log data sets for future NFS usage. Note that the last data buffer is lost when the log is switched.

GFSA931I (*procname*) NETWORK FILE SYSTEM SERVER LOGGING IS TERMINATED.

Explanation: The NFS logging is ended. This can be caused by a "no space" condition of the log data set.

In the message text:

procname The name of the start procedure.

System action: NFS processing continues.

System programmer response: Allocate a larger log data set for future NFS usage.

GFSA932I (*procname*) LOG DATA SET *text* IS FLUSHED.

Explanation: The data buffer of the active log data set is flushed to disk. If the log is written to standard error, *text* is STDERR. Otherwise, *text* is the associated data definition of the active log data set.

In the message text:

procname The name of the start procedure.

System action: NFS processing continues.

GFSA933I (*procname*) LOG DATA SET IS RE-INITIALIZED.

Explanation: The log data set is reinitialized.

In the message text:

procname The name of the start procedure.

System action: NFS processing continues.

GFSA934I (*procname*) NFSLOG1 OR NFSLOG2 DD STATEMENTS NOT DEFINED.

Explanation: The NFSLOG1 and NFSLOG2 DD statements are not coded in the NFS startup cataloged procedure.

In the message text:

procname The name of the start procedure.

System action: The server stops.

System programmer response: Code the NFSLOG1 and NFSLOG2 DD statements and allocate the associated log data sets.

GFSA935I (*procname*) SWITCHED-TO LOG IN ERROR. NETWORK FILE SYSTEM SERVER LOGGING CONTINUED ON DD:*text*.

Explanation: An operator issued a **switchlog** command, but an open error is detected in the new log data set. NFS logging is continued in the original log data set.

In the message text:

procname
The name of the start procedure.

text The value of *text* is the data definition associated with the original log data set.

System action: NFS processing continues.

Operator response: Fix the nonactive log data set.

GFSA936I (*procname*) NETWORK FILE SYSTEM SERVER LOG *text* SET TO FORCELOG.

Explanation: The NFS server log was closed to force all log data to disk immediately.

In the message text:

procname The name of the start procedure.

System action: NFS continues.

GFSA937E MODIFY MAPFILE COMMAND – ERROR IN SIDEFILE DATA SET – – SIDEFILE IS NOT REFRESHED.

Explanation: This message is the response from the **mapfile** operand of the **modify** command, indicating that one or more errors were detected in the side files data set (for example, the side files data set cannot be opened).

System action: NFS processing continues. The existing side files data in storage is not changed.

Operator response: Notify your system programmer.

System programmer response: Review previous error messages for detailed information as to the specific error in the side files data set.

GFSA938I SIDE FILE WAS NOT LOADED BEFORE.

Explanation: This message is the response from the **mapfile** of the **modify** command, indicating that the command was ignored because the side file was not loaded.

System action: NFS processing continues with the side files data unchanged.

GFSA939E (*procname*) Module *text1* release *text2* in *text3* is incompatible with Network File System Server release *text4*

Explanation: During the Network File System server initialization a module was found in a task which is at an incompatible release level.

In the message text:

Text1 Module which is in error.

Text2 Release level of the module in error.

Text3 Task in which the module in error was found.

Text4 Network File System server release level.

System action: The Network File System server terminates.

Operator response: Notify the system programmer.

System programmer response: Verify that the last maintenance was installed correctly by SMP/E. When a maintenance update is installed for a module, SMP should install it in all load modules.

Source: z/OS NFS Server

GFSA940E (*procname*) Module *text1* is incompatible between Network File System Server load modules *text2 (text3)* and *text4 (text5)*

Explanation: During the Network File System server initialization a module was found to be at different maintenance levels in two load modules of the server.

In the message text:

Text1 Module which is at inconsistent maintenance levels.

Text2 Name of first load module containing the module in error.

Text3 APAR maintenance version of the module in load module *text2*.

Text4 Name of second load module containing the module in error.

Text5 APAR maintenance version of the module in the other load module *text4*.

System action: The Network File System server terminates.

Operator response: Notify the system programmer.

System programmer response: Verify that the last maintenance was installed correctly by SMP/E. When a maintenance update is installed for a module, SMP should install it in all load modules.

Source: z/OS NFS Server

GFSA941I UNMNTXXX: MODIFY COMMAND WAS COMPLETED SUCCESSFULLY.

Explanation: This message is the response from the **unmntxxx** operand of the **modify** command, indicating the command completed normally.

System action: The NFS server immediately unmounts the specified type of mount points.

GFSA942I UNMNTXXX: MODIFY COMMAND IGNORED – MOUNT POINT *mount_point* **IS IN USE BY USER** *user_name*

Explanation: This message is the response from the **unmnt***xxx* operand of the **modify** command, indicating that *mount_point* was ignored because the mount point is in use by another user, *user_name*.

System action: The NFS server processing continues.

Operator response: Reenter the same operand of the **modify** command and try again later.

GFSA943E (*procname***) EXPORTS: UNEXPECTED OPTION (***text***) WAS DETECTED IN THE EXPORTS FILE.**

Explanation: The option information provided in *text* is incorrect. This error could occur as a result of unexpected blanks, incorrect syntax, or mutually exclusive options (for example, both **ro** and **rw**). This message might occur after performing the EXPORTFS command.

In the message text:

procname The name of the start procedure.

System action: The NFS server continues processing.

Operator response: Reenter the same operand of the **modify** command.

GFSA944I (*procname***)** *text1* **Network File System Server release** *text2* **last APAR** *text3* **last changed module:** *text4* **compiled at** *text5*

Explanation: This message is reported by the Network File System server in response to the MODIFY mvsnfs,VERSION operator command. It reports the current maintenance level of the server.

In the message text:

Text1 System level (z/OS, for example).

Text2 Network File System server release level.

Text3 Last APAR maintenance version of Network File System server.

Text4 Name of one module installed by the last maintenance level.

Text5 Compilation date/time of the module.

System action: The Network File System server continues.

Operator response: None.

System programmer response: None.

Source: z/OS NFS Server

GFSA945I (*procname***)** *text1* | *text2* | *text3*

Explanation: This message is reported by the Network File System server in response to the operator command MODIFY mvsnfs,VERSION=ALL or MODIFY mvsnfs,VERSION=MODULE. Message GFSA945I is generated for each module.

In the message text:

Text1 Module name.

Text2 APAR maintenance level of the module.

Text3 Compile date/time of the module.

System action: The Network File System server continues.

Operator response: None.

System programmer response: None.

Source: z/OS NFS Server

GFSA946I (*procname***) Total modules:** *digit*

Explanation: This message is returned by the Network File System server in response to the MODIFY mvsnfs,VERSION=ALL operator command, after message GFSA945I lists the information on all the modules. This message reports the total module count *digit*.

System action: The Network File System server continues.

Operator response: None.

System programmer response: None.

Source: z/OS NFS Server

GFSA947I (*procname***)** *text1* | *text2* | *text3*

Explanation: This message is reported by the Network File System server during initialization if debug tracing is turned on. Message GFSA947I is generated for each module.

In the message text:

Text1 Module name.

Text2 APAR maintenance level of the module.

Text3 Compile date/time of the module.

System action: The Network File System server continues.

Operator response: None.

System programmer response: None.

Source: z/OS NFS Server

GFSA948I (*procname*) Non-existent module: *text1*

Explanation: This message is returned by the Network File System server in response to the MODIFY mvsnfs,VERSION=*text1* operator command, indicating that the module name specified as *text1* does not exist.

System action: The Network File System server continues.

Operator response: None.

System programmer response: None.

Source: z/OS NFS Server

GFSA949I COMMAND *text* NOT VALID.

Explanation: The operator entered the **operator** command or the operand *text*, which is not valid.

System action: NFS processing continues.

Operator response: Reenter the operand of the **modify** command with the correct syntax.

GFSA950I Unknown flag '-*character*'

Explanation: The character *character* specified in the **mvslogin**, **mvslogout**, or **showattr** command is not a valid option. A usage message might follow this message.

User response: See *z/OS Network File System Guide and Reference* for a description of the valid options used with the command.

GFSA951I *text* : can't find name for uid *d_digits*.

Explanation: There was an error reading information for the user ID *d-digits* from the etc/passwd file.

User response: Correct the etc/passwd file and retry the command.

GFSA952I Retyped password does not match

Explanation: The password entered when message GFSA975I displayed does not match the password entered when message GFSA974A displayed.

User response: Restart the **mvslogin** command sequence.

GFSA953I Password change required by host.

Explanation: The multiple virtual system (MVS) password for the user ID passed to the host expired. A new password is required. Message GFSA974I follows this message.

GFSA954I Host *text1* returned error *d_digits*: *text2*

Explanation: An error was detected during **mvslogin** processing. Host *text1* returned error code *d_digits* and message *text2* to the client.

User response: The password or user ID might be incorrect. Restart the **mvslogin** command sequence and use the correct password or user ID.

GFSA955I *text* logged in ok.

Explanation: The MVS user ID *text* was logged in without any errors.

GFSA956I usage: *text* [-pn][-g group][-a account] hostname [mvs_username]

Explanation: This message provides usage information for the **text** command.

User response: Enter the command using the correct syntax.

GFSA957I Host *text1* returned error *d_digits*: *text2*

Explanation: An error was detected during **mvslogout** processing. Host *text1* returned error code *d_digits* and message *text2* to the client.

User response: Notify your system programmer.

GFSA958I uid *text* logged out ok.

Explanation: The user ID *text* logged out successfully.

GFSA959I usage: *text* hostname

Explanation: This message provides usage information for the *text* command.

User response: Enter the command using the correct syntax.

GFSA960I *text1*: host "*text2*" unknown.

Explanation: The host *text2*, specified in the command *text1*, is not known to the network.

User response: Correct the host name specified and retry the command.

GFSA961I *text1*: *text2*

Explanation: The command *text1* received an error when trying to create a client transport handle using the **clntudp_create** TCP/IP remote procedure call. The value of *text2* is the message produced by the **clnt_pcreateerror** TCP/IP procedure. This message is issued in any of these cases:

- The host name is unknown.
- The host is not operational.
- The NFS on the named host is not operational.

User response: The user response is one of these actions:

- Correct the specified host name and retry the command.
- Make sure the specified host is operational and retry the command.
- Make sure the NFS server on the named host is operational and retry the command.

GFSA962E **Unable to create CLNT handle. rc =** *d_digit1***, rsn =** *hex_digit1***.**

Explanation: The NFS client was unable to create a handle for communicating with the NFS server. The request failed with return code *d_digit1* and reason code *hex_digit1*.

System action: The Network File System Client request fails.

Operator response: None.

System programmer response: Verify that the network connection to the NFS Server is operational and that the NFS Server is functional.

GFSA963E **Call to remote Network File System Server failed. RV=***hex_digit1* **RC=***hex_digit2* **RSN=***hex_digit3***.**

Explanation: The request to the Network File System Server failed. The request failed with return value *hex_digit1*, return code *hex_digit2* , and reason code *hex_digit3*.

System action: The Network File System Client request fails.

Operator response: None.

System programmer response: Verify that the network connection to the NFS Server is operational and that the NFS Server is functional.

GFSA964I *text***: Error: cannot determine server.**

Explanation: The *text* command found the mount path, but the server name was not returned by the local operating system service that keeps mount point information.

User response: Correct the mount point table and retry the command.

GFSA965I *text1***: Error:** *text2* **mounted from server** *text3***, not** *text4***.**

Explanation: The wrong host name was specified for the *text1* command. *text2* is mounted from server *text3* instead of server *text4*.

User response: Specify the command again with the correct host name.

GFSA966I *text***: Error: unknown return from usage routine.**

Explanation: The usage routine used for the **text** command returned an unknown error code.

User response: Contact your programming support personnel.

GFSA967I **Host Error:** *text***.**

Explanation: The host returned an error and the message *text*. The error might be due to either of these causes:

- A porting failure occurred for the **showattr** command.
- The NFS is not compatible with the **showattr** command on the client.

User response: Contact your programming support personnel.

GFSA968I **Error: Drive** *text* **not mounted.**

Explanation: The drive *text* was not mounted. The **showattr** command cannot show the attributes for this drive.

User response: Mount the drive and reissue the command.

GFSA969I **Error: Can't open** *text* **for read.**

Explanation: The file *text* could not be opened to read the mount path.

User response: Correct the file and reissue the command.

GFSA970I **Error: Directory** *text* **not mounted.**

Explanation: The directory *text* was not mounted. The **showattr** command cannot show the attributes for this directory.

User response: Mount the directory and reissue the command.

GFSA971I Error: filesystem *text* is local.

Explanation: The file system *text* is not a NFS file system. The **showattr** command is for NFS file systems only.

User response: Reissue the command for a NFS file system.

GFSA972I usage: *text1* [-t] hostname [*text2*]

Explanation: This message provides usage information for the *text1* command.

In the message text:

text2 The value of *text2* is the operating-system-dependent mount-point format.

GFSA973A Enter MVS password:

Explanation: The NFS requires a password for the user.

User response: If a user ID was specified in the **mvslogin** command, enter the password for that user ID. If no user ID was specified, the name from etc/passwd for the user ID that issued the **mvslogin** command was passed to the NFS. Enter the password for this user.

GFSA974A Enter new MVS password:

Explanation: The password for the user ID passed to NFS has expired.

User response: Enter a new password.

GFSA975A Retype new MVS password:

Explanation: The system requires the new password to be entered twice for verification.

User response: Reenter the new password.

GFSA976I *text1*: *text2*

Explanation: The command *text1* received an error when trying to create a client transport handle using the **clnt_call** TCP/IP remote procedure call.

In the message text:

text1 The value of *text2* is the message produced by the **clnt_perror** TCP/IP procedure.

User response: Contact your system administrator.

GFSA977I *text*:

Explanation: The command *text* received an error when trying to create a client transport handle using the **clnt_call** TCP/IP remote procedure call (RPC). This message is followed by the message produced by the **clnt_perror** TCP/IP procedure.

User response: Contact the system administrator.

GFSA978I *text* logged in ok. Mismatch in uid/gid: z/OS UNIX uid is *digit_1*, gid is *digit_2*, client uid is *digit_3*, gid is *digit_4*

Explanation: The z/OS UNIX user ID or group ID does not match the client machine user ID or group ID. The authentication is successful and the message is for information only.

GFSA979E Unable to create UDP socket rc = *d_digit1*.

Explanation: The NFS client was unable to create a UDP protocol socket for communicating with the NFS server. The request failed with return code *d_digit1*.

System action: The Network File System Client request fails.

Operator response: None.

System programmer response: Verify that the network connection to the NFS Server is operational and that the NFS Server is functional.

GFSA980E Unable to get the program's port.

Explanation: The NFS client was unable to find the port for communicating with the NFS server.

System action: The Network File System Client request fails.

Operator response: None.

System programmer response: Verify that the network connection to the NFS Server is operational and that the NFS Server is functional.

GFSA981I Unable to establish TCP connection rc = *d_digit1*. Trying UDP ...

Explanation: The NFS client was unable to establish a TCP protocol connection to the NFS server. The request failed with return code *d_digit1*.

System action: The Network File System Client will attempt to create a UDP protocol connection.

Operator response: None.

System programmer response: None.

GFSA982E The program *d_digit1* is not registered on *string*.

Explanation: The requested program *d_digit1* is not registered on host *string*.

System action: The Network File System Client request fails.

Operator response: None.

System programmer response: Verify that the network connection to the NFS Server is operational and that the NFS Server is functional.

GFSA983E Remote host does not have *string* interface.

Explanation: The remote host does not support the desired IP interface *string*. The possible values of *string* are:

AF_INET
Internet Protocol version 4

AF_INET6
Internet Protocol version 6

System action: The Network File System Client request fails.

Operator response: None.

System programmer response: Verify that the requested IP version is operational.

GFSA984E Error in authunix_create.

Explanation: The Network File System Client was unable to create an Auth_sys security connection to the NFS server.

System action: The Network File System Client request fails.

Operator response: None.

System programmer response: Verify that the system is properly configured.

GFSA985E Cannot get local host name

Explanation: The Network File System client utility was unable to determine the NFS Server host name for the specified mount point.

System action: The Network File System Client utility request fails.

Operator response: None.

System programmer response: None.

GFSA986I The Network File System client utility was unable to get the attribute information for the specified mount point.

Explanation: The Network File System client utility was unable to determine the NFS Server host name for the specified mount point.

System action: The Network File System Client utility request fails.

Operator response: None.

System programmer response: None.

GFSA987I Directory *string* is not a supported Network File System type of directory.

Explanation: The specified directory path *string* is not a Network File System mounted directory.

System action: The Network File System Client request fails.

User response: Correct the directory path specification.

GFSA990I *text*

Explanation: The LOGON or FILE SECURITY user exit routine returned the message *text* to the NFS server.

System action: The Network File System server processing continues.

User response: No action required.

GFSA991E (*procname*) MESSAGE FORMAT FROM USER EXIT ROUTINE(S) IS INCORRECT. USER *text* EXIT ROUTINE(S) HAS(HAVE) ENDED.

Explanation: The *test* exit routine returned a message with an incorrect length.

In the message text:

procname
The name of the start procedure.

text The value of *text* can be either `LOGON AND FILE SECURITY` or `FILE SECURITY`.

System action: The NFS server processing continues without the user exit routine(s).

Operator response: Record the operator console message and notify the NFS system programmer.

System programmer response: Correct the user exit routine(s), relink the user exit routine(s), and restart the NFS server.

GFSA996E (*procname*) INSTALLATION DEFAULT TRANSLATION TABLE CANNOT BE INITIALIZED.

Explanation: The NFS ends because one of these conditions happens during NFS startup:

- The NFSXLAT DD statement is not coded in the NFS startup catalog procedure, and the **xlat** processing attribute is specified.
- The translation table data set defined in the NFSXLAT DD statement is not a PDS or PDSE.
- The translation table specified in the translation table data set cannot be found.
- The translation table contained in the translation table data set is in an incorrect format.

In the message text:

procname The name of the start procedure.

System action: NFS processing stops.

System programmer response: Take the following actions before restarting the NFS:

- Code the NFSXLAT DD statement in the NFS startup catalog procedure correctly.
- Allocate the translation table data set defined in the NFSXLAT DD statement.
- Make sure the translation table in the translation table data set exists.
- Make sure the translation table specified in the translation table data set is in the correct format.

GFSA997I ERROR IN READING TRANSLATION TABLE, *text***.**

Explanation: The NFS tried to read the translation table, *text*, during the mount but was unsuccessful. The value of *text* is the name of the member that contains the translation table. The reason might be one of these causes:

- The NFSXLAT DD statement is not coded in the NFS startup catalog procedure.
- The NFSXLAT DD statement is not coded correctly in the NFS startup catalog procedure.
- The translation table data set defined in the NFSXLAT DD statement is not a PDS or PDSE.
- The translation table specified in the translation table data set cannot be found.
- The format of the translation table contained in the translation table data set is not valid.

System action: NFS processing stops.

User response: This error occurs during the mount operation. Contact your system programmer for the correct member name for the translation table.

GFSA998I TRANSLATION TABLE *text* **IS LOADED.**

Explanation: The NFS loaded the translation table, *text*, successfully.

System action: NFS processing continues.

Client messages

This is a listing of the messages that the NFS client generates. A message, explanation, and recommended action are supplied where applicable. Data is substituted for any part of a message shown here in *italics*.

Messages that appear on the z/OS operator's console for the NFS client will be in the following example format: **GFSC700I z/OS NETWORK FILE SYSTEM CLIENT (HDZ11VC) STARTED**

Table 58 shows the message format on the NFS client operators console.

Table 58. NFS client z/OS operators console message format

GFSC	Component identifier for the NFS client	
700	A unique message number	
I	Message type:	
	A	Action; the user must perform a specific action.
	E	Eventual action; the user must perform an action when time is available.
	I	Informational; no user action is required.
z/OS NETWORK FILE SYSTEM CLIENT (HDZ11VC) STARTED	Message text	

Messages appear in the NFS client log data set in the same format as in the following example:

12:34:18 GFSC100E (E) CCXDR 11 XDR_DISP: RPC REQUEST (7F638E30) FAILED, RETURN VALUE -1 RETURN CODE 00000467 REASON CODE 5 (TIMED OUT)

Table 59 on page 281 shows the message format in the NFS client log data set.

Table 59. NFS client log data set message format

12:34:18	The time stamp (*hours*:*minutes*:*seconds*)	
GFSC	Component identifier for the NFS client	
100	A unique message number	
E	Message type:	
	A	Action; the user must perform a specific action.
	E	Eventual action; the user must perform an action when time is available.
	I	Informational; no user action is required.
(E)	The message level: E (error), W (warning), or I (informational)	
CCXDR 11	'CCXDR' is the last 5 characters of CSECT name. '11' is the message sequence within the function.	
XDR_DISP	The first 8 characters of the function name	

Table 59. NFS client log data set message format (continued)

RPC REQUEST (7F638E30) FAILED, RETURN VALUE -1 RETURN CODE 00000467 REASON CODE 5 (TIMED OUT)	Message text

The messages are listed in numerical order (the time stamp, message level, and programming support information are not shown).

Notes:

1. Messages GFSC098I and GFSC099I are intended for IBM support personnel to use when they are performing diagnosis. They do not indicate a problem with NFS but do provide statistics on NFS processing. As such, an extensive number of GFSC098I and GFSC099I messages may be issued.

Table 60 shows common variables in the message text.

Table 60. Common variables

Variable	Meaning
retv	Decimal return value
retc	Decimal return code
rsnc	Decimal reason code
returncd	8-digit hexadecimal return code – see Chapter 20, "Return codes," on page 299
reasoncd	8-digit hexadecimal reason code – see Chapter 21, "Reason codes," on page 309
h_digit	8-digit hexadecimal address
d_digit	Decimal digits
dsname	Data set name
text	Place holder for long text of different lengths

GFSC100E RPC REQUEST (*h_digit*) FAILED, RETURN VALUE -1 RETURN CODE *returncd* REASON CODE *rsnc* (*text*)

Explanation: A remote procedure call (RPC) request failed.

In the message text:

h_digit	The value of *h_digit* is the request block address.
rsnc	The reason code, *rsnc*, is the return code returned from TCP/IP.
text	The value of *text* is the failure reason.
returncd	For the explanation of return code *returncd*; see *z/OS UNIX System Services Messages and Codes*.

System action: The client processing continues.

User response: See the return code *returncd* in *z/OS UNIX System Services Messages and Codes* and reason code *rsnc* in the TCP/IP message manual.

GFSC101E NETWORK FILE SYSTEM SERVER REQUEST FAILED (*h_digit*), z/OS UNIX RETURN CODE *returncd* NETWORK FILE SYSTEM SERVER RETURN CODE *retc* (*text*)

Explanation: The NFS server failed the request from the client.

In the message text:

h_digit	The value of *h_digit* is the request block address.
retc	The return code, *retc*, is returned from the NFS server.
text	The value of *text* is the failure reason for the request.
returncd	For the explanation of return code *returncd*, see *z/OS UNIX System Services Messages and Codes*.

System action: The client processing continues.

User response: See the return code *returncd* in *z/OS*

UNIX System Services Messages and Codes and the return code *retc* in the *Network File System Protocol Specification, RFC 1094* documentation.

GFSC102E RPC REQUEST (*h_digit*) FAILED, RETURN VALUE -1 RETURN CODE *returncd* REASON CODE *rsnc*

Explanation: A remote procedure call (RPC) request failed.

In the message text:

h_digit The value of *h_digit* is the request block address.

rsnc The reason code, *rsnc*, is the return code returned from TCP/IP.

returncd The explanation of return code *returncd* is in *z/OS UNIX System Services Messages and Codes*.

System action: The client processing continues.

User response: See the return code *returncd* in *z/OS UNIX System Services Messages and Codes* and the reason code *rsnc* in the TCP/IP message document.

GFSC103E NETWORK FILE SYSTEM SERVER REQUEST FAILED (*h_digit*), z/OS UNIX RETURN CODE *returncd* NETWORK FILE SYSTEM SERVER RETURN CODE *retc*

Explanation: The NFS server failed the request from the client.

In the message text:

h_digit The value of *h_digit* is the request block address.

retc The return code, *retc*, is returned from the NFS server.

returncd For the explanation of return code *returncd* see *z/OS UNIX System Services Messages and Codes*.

System action: The client processing continues.

User response: See the return code *returncd* in *z/OS UNIX System Services Messages and Codes*and the return code *retc* in the *Network File System Protocol Specification, RFC 1094* documentation.

GFSC105E READ FAILED, RETURN VALUE -1 RETURN CODE *returncd* REASON CODE *reasoncd*.

Explanation: The system detected an error, *returncd*, while reading a block of data from a remote file.

System action: The read operation ends. NFS client processing continues.

User response: See *z/OS UNIX System Services Messages and Codes* for a description of the return code *returncd* and the reason code *reasoncd* to determine corrective action.

GFSC106E WRITE FAILED, RETURN VALUE -1 RETURN CODE *returncd* REASON CODE *reasoncd*.

Explanation: The system detected an error, *returncd*, while writing a block of data to a remote file.

System action: The write operation ends. NFS client processing continues. The remote file might not be complete.

User response: See *z/OS UNIX System Services Messages and Codes* for a description of the return code *returncd* and the reason code *reasoncd* to determine corrective action.

GFSC107E FLUSH FAILED, RETURN VALUE -1 RETURN CODE *returncd* REASON CODE *reasoncd*.

Explanation: The system detected an error, *returncd*, while flushing cached data to a remote file during close processing.

System action: The write operation ends. NFS client processing continues. The remote file might not be complete.

User response: See *z/OS UNIX System Services Messages and Codes* for a description of the return code *returncd* and the reason code *reasoncd* to determine corrective action.

GFSC110E *text* FAILED, RETURN VALUE -1 RETURN CODE *returncd* REASON CODE *reasoncd*

Explanation: The NFS client has detected an error in the function *text*.

System action: The request ended. NFS client processing continues.

User response: See *z/OS UNIX System Services Messages and Codes* for a description of the return code *returncd* and the reason code *reasoncd* to determine corrective action.

System programmer response: Collect the detail trace log from the client and from the server, if any.

GFSC200E VFS_MOUNT FAILED, RETURN VALUE -1 RETURN CODE *returncd* REASON CODE *reasoncd*.

Explanation: The **mount** command failed because of error *returncd*.

System action: The **mount** command ended

abnormally. NFS client processing continues.

User response: See *z/OS UNIX System Services Messages and Codes* for a description of the return code, *returncd* and the reason code *reasoncd* to determine the corrective action, and reissue the **mount** command.

GFSC201E NETWORK FILE SYSTEM CLIENT DOES NOT SUPPORT SYNCHRONOUS MOUNT REQUEST.

Explanation: The NFS client supports only an asynchronous mounts.

System action: The **mount** command ended with an error. NFS client processing continues.

User response: Reissue the **mount** command with the **asynchronous** option.

GFSC202E A FILE SYSTEM WITH THE SAME NAME IS ALREADY MOUNTED.

Explanation: The system cannot mount on an existing mount point.

System action: The **mount** command ended with an error. No mount point was established. NFS client processing continues.

User response: Reissue the **mount** command with a different mount point.

GFSC203E PARSING MOUNT OPTION FAILED, RETURN VALUE -1 RETURN CODE *returncd* **REASON CODE** *reasoncd* **OPTION=***'text'***.**

Explanation: The mount option *text* was incorrectly specified.

System action: The **mount** command ended with an error. No mount point is established. NFS client processing continues.

User response: See *z/OS UNIX System Services Messages and Codes* for a description of the return code *returncd* and see Table 67 on page 309 for more details of the reason code *reasoncd*. Correct the mount option, and reissue the **mount** command.

GFSC204E VFS_UMOUNT FAILED, RETURN VALUE -1 RETURN CODE *returncd* **REASON CODE** *reasoncd***.**

Explanation: The **umount** or **unmount** command failed.

System action: The **umount** or **unmount** command ended with an error. The mount point might still exist. NFS client processing continues.

User response: See *z/OS UNIX System Services Messages and Codes* for a description of the return code *returncd* and the reason code *reasoncd* to determine the corrective action. Correct it and reissue the **unmount** command.

GFSC205E VFS_STATFS FAILED, RETURN VALUE -1 RETURN CODE *returncd* **REASON CODE** *reasoncd***.**

Explanation: The VFS_STATFS operation failed. The system detected an error, *returncd*, while trying to get the status of a remote file system.

System action: The VFS_STATFS operation ended with an error. NFS client processing continues.

User response: See *z/OS UNIX System Services Messages and Codes* for a description of the return code *returncd* and the reason code *reasoncd* to determine the corrective action.

GFSC206E VFS_SYNC FAILED, RETURN VALUE -1 RETURN CODE *returncd* **REASON CODE** *reasoncd***.**

Explanation: The VFS_SYNC operation failed. The system detected an error, *returncd* while flushing cached data of remote files.

System action: The VFS_SYNC operation ended with an error. The remote files might not be complete. NFS client processing continues.

User response: See *z/OS UNIX System Services Messages and Codes* for a description of the return code *returncd* and the reason code *reasoncd* to determine the corrective action.

GFSC207E VFS_RECOVER FAILED, RETURN VALUE -1 RETURN CODE *returncd* **REASON CODE** *reasoncd***.**

Explanation: The VFS_RECOVER operation failed. The system detected an error, *returncd*, while trying to recover from a previous abend.

System action: The VFS_RECOVER operation ended with an error. NFS client processing continues.

User response: See *z/OS UNIX System Services Messages and Codes* for a description of the return code *returncd* and the reason code *reasoncd* to determine the corrective action. If the error occurs and the program is not in error, see the messages in the client log data sets for more information. Search problem-reporting databases for a fix for the problem. If no fix exists, contact the IBM Support Center. Provide all the printed output and copies of output data sets related to the problem.

GFSC208E VFS_PFSCTL FAILED, RETURN VALUE -1 RETURN CODE *returncd* **REASON CODE** *reasoncd***.**

Explanation: The VFS_PFSCTL operation failed.

System action: The VFS_PFSCTL operation ended with an error. NFS client processing continues.

User response: See *z/OS UNIX System Services Messages and Codes* for a description of the return code *returncd* and the reason code *reasoncd* to determine the corrective action. If the error occurs and the program is not in error, look at the messages in the client log data sets for more information. Search problem-reporting databases for a fix for the problem. If no fix exists, contact the IBM Support Center. Provide all the printed output and copies of output data sets related to the problem.

GFSC209E THE IP ADDRESS OF REMOTE HOST NAME, *hostname*, COULD NOT BE RESOLVED.

Explanation: The VFS_MOUNT operation failed. The mount processing failed when it tried to resolve the remote host name, *hostname*, to a dotted IP address.

System action: The **mount** command ended with an error. NFS client processing continues.

User response: Correct the remote host name, or use the dotted IP address of the remote host, and reissue the **mount** command.

GFSC210E NFS SERVER *hostname* DOES NOT SUPPORT NFS VERSION 3 PROTOCOL WHILE 'VERS(3) WAS SPECIFIED

Explanation: The VFS_MOUNT operation failed. The mount processing failed because the server does not support NFS version 3 protocol, which the user requested with the **mount** parameter `vers(3)`.

System action: The **mount** command ended with an error. NFS client processing continues.

User response: Verify that the server actually does not support NFS version 3 protocol. Remove `vers(3)` from the **mount** parameter, and reissue the **mount** command.

GFSC211E NFS SERVER *hostname* DOES NOT SUPPORT 'AUTH_SYS' AUTHENTICATION.

Explanation: The VFS_MOUNT operation failed. The **mount** processing failed because the server does not do AUTH_SYS authentification. z/OS does not support other authentification, such as Kerberos.

System action: The **mount** command ended with an error.

User response: Verify that the server actually does not support AUTH_SYS authentification. Notify the server system administrator.

GFSC212E MOUNT FAILED BECAUSE OF CONVERSION SERVICE CONNECTION FAILURE CCSID *ccsid* RETURN CODE *returncd* REASON CODE *reasoncd*

Explanation: Coded character set identifiers (CCSID) specified in the **mount** command are not supported by the conversion service.

System action: The **mount** command ended abnormally. NFS client processing continues.

User response: See *z/OS UNIX System Services Messages and Codes* for more information about the return code *returncd*, and the reason code *reasoncd* to determine corrective action.

System programmer response: Check the availability of the specified CCSIDs.

GFSC213E REQUEST (*requestid*) : THE NFS SERVER *hostname* DOES NOT SUPPORT THE SPECIFIED 'VERS' AND/OR 'PROTO'.

Explanation: The **mount** request *requestid* failed because the user-specified **vers** and **proto** was not supported by the NFS server on *hostname*.

System action: The **mount** command ended with an error. NFS client continues processing.

User response: Either let the NFS client choose the compatible **vers** and **proto**, or determine the NFS server capabilities (by **orpcinfo**) and reissue the **mount** command with the correct **vers** or **proto**.

GFSC214E REQUEST (*requestid*) : THE NFS SERVER *hostname* DOES NOT HAVE NFS REGISTERED ON PORT 2049

Explanation: The request *requestid* failed because the NFS server on *hostname* did not register or use port 2049.

System action: The operation ended with an error. NFS client continues processing.

User response: Use **orpcinfo** to verify the server and correct the server.

GFSC300E MISSING LEFT PARENTHESIS IN *text* KEYWORD.

Explanation: The specified keyword, *text*, is missing a left parenthesis.

System action: NFS client processing stops if the error occurs in the NFS client installation parameter. The **mount** command failed if the error is in the **mount** parameter.

Operator response: Record the z/OS operator console message and notify the system programmer.

System programmer response: If the error is in the **mount** parameter, correct the parameter and reissue the **mount** command. If the error occurs in the NFS client installation parameter, correct the parameter, stop z/OS UNIX, and restart z/OS UNIX.

GFSC301E PARSE FAILED ON NUMERIC FIELD FOR *text* KEYWORD.

Explanation: The specified keyword, *text*, contains alphabetic data in a numeric field.

System action: NFS client processing stops if the error is in the NFS client installation parameter. The **mount** command failed if the error is in the **mount** parameter.

Operator response: Record the z/OS operator console message and notify the system programmer.

System programmer response: If the error is in the **mount** parameter, correct the parameter and reissue the **mount** command. If the error occurs in the NFS client installation parameter, correct the parameter, stop z/OS UNIX, and restart z/OS UNIX.

GFSC302E MISSING RIGHT PARENTHESIS IN *text* KEYWORD.

Explanation: The specified keyword, *text*, is missing a right parenthesis.

System action: NFS client processing stops if the error is in the NFS client installation parameter. The **mount** command failed if the error is in the **mount** parameter.

Operator response: Record the z/OS operator console message and notify the system programmer.

System programmer response: If the error is in the **mount** parameter, correct the parameter and reissue the **mount** command. If the error occurs in the NFS client installation parameter, correct the parameter, stop z/OS UNIX, and restart z/OS UNIX.

GFSC303E THE SPECIFIED VALUE *d_digit1* IS NOT IN THE RANGE OF *d_digit2* TO *d_digit3* FOR *text* KEYWORD.

Explanation: The value *d_digit1* specified in the keyword *text* must be between the minimum value, *d_digit2*, and the maximum value, *d_digit3*.

System action: NFS client processing stops if the error is in the NFS client installation parameter. The **mount** command failed if the error is in the **mount** parameter.

Operator response: Record the z/OS operator console message and notify the system programmer.

System programmer response: If the error is in the **mount** parameter, correct the parameter and reissue the **mount** command. If the error occurs in the NFS client installation parameter, correct the parameter, stop z/OS UNIX, and restart z/OS UNIX.

GFSC304E PARSE FAILED ON ALPHABETIC FIELD FOR *text* KEYWORD.

Explanation: The specified keyword, *text*, contains numeric data for an alphabetic field.

System action: NFS client processing stops if the error is in the NFS client installation parameter. The **mount** command failed if the error is in the **mount** parameter.

Operator response: Record the z/OS operator console message and notify the system programmer.

System programmer response: If the error is in the **mount** parameter, correct the parameter and reissue the **mount** command. If the error occurs in the NFS client installation parameter, correct the parameter, stop z/OS UNIX, and restart z/OS UNIX.

GFSC305E INCORRECT OPTION *text1* SPECIFIED FOR *text2* KEYWORD, VALID OPTION IS Y OR N.

Explanation: An incorrect option, *text1*, was specified for the keyword *text2*.

System action: NFS client processing stops if the error is in the NFS client installation parameter. The **mount** command failed if the error is in the **mount** parameter.

Operator response: Record the z/OS operator console message and notify the system programmer.

System programmer response: If the error is in the **mount** parameter, correct the option for the keyword *text2* and reissue the **mount** command. If the error occurs in the NFS client installation parameter, correct the option for the keyword *text2*, stop z/OS UNIX, and restart z/OS UNIX.

GFSC307E *text* IS AN INCORRECT KEYWORD FOR MOUNT PARAMETER.

Explanation: The value of *text* can be specified only as an installation parameter.

System action: The **mount** command failed.

System programmer response: Correct the **mount** parameter keyword *text*.

GFSC308E REQUEST (*requestid*) : XLAT(Y) CANNOT BE SPECIFIED AS A MOUNT PARAMETER WHEN TAG OPTION IS ALSO SPECIFIED.

Explanation: The **mount** request with the *requestid* code failed because the user specified both the **xlat(Y)** and the **tag** options together. This is not allowed.

System action: The **mount** command ended with an error. NFS client continues processing.

User response: Either specify **xlat(Y)** to have the NFS client do text translation based on the **cln_ccsid** and

srv_ccsid values, or specify the **tag** option with the correct coded character set identifier (CCSID) to have the translation done by the LFS based on the CCSID in the **tag** option.

GFSC309E UNKNOWN KEYWORD ENCOUNTERED AROUND POSITION *d_digit***.**

Explanation: The keyword specified in position *d_digit* is not a valid keyword.

System action: NFS client processing stops if the error is in the NFS client installation parameter. The **mount** command failed if the error is in the **mount** parameter.

Operator response: Record the z/OS operator console message and notify the system programmer.

System programmer response: If the error is in the **mount** parameter, correct the keyword and reissue the **mount** command. If the error occurs in the NFS client installation parameter, correct the keyword, stop z/OS UNIX, and restart z/OS UNIX.

GFSC310I READAHEAD AND DELAYWRITE OPTIONS WILL BE IGNORED AS DATACACHING IS OFF.

Explanation: The keywords **readahead** and **delaywrite** are ignored because **datacaching** has been set to OFF.

System action: NFS client processing continues.

Operator response: Record the z/OS operator console message and notify the system programmer.

System programmer response: Check the parameters to make sure that **datacaching** should be OFF.

GFSC311I CLN_CCSID AND SRV_CCSID WILL BE IGNORED AS XLAT OPTION IS OFF.

Explanation: The keywords **cln_ccsid** and **srv_ccsid** are ignored because **xlat** has been set to OFF.

System action: NFS client processing continues.

Operator response: Record the z/OS operator console message and notify the system programmer.

System programmer response: Check the parameters to make sure that **xlat** should be OFF.

GFSC312I ACREGMIN, ACREGMAX, ACDIRMIN, AND ACDIRMAX OPTIONS WILL BE IGNORED AS ATTRCACHING IS OFF.

Explanation: The keywords **acregmin**, **acregmax**, **acdirmin**, and **acdirmax** are ignored because **attrcaching** has been set to OFF.

System action: NFS client processing continues.

Operator response: Record the z/OS operator console message and notify the system programmer.

System programmer response: Check the parameters to make sure that **attrcaching** should be OFF.

GFSC313I RETRANS OPTION WILL BE IGNORED AS HARD OPTION IS ON.

Explanation: The keyword **retrans** is ignored because **hard** has been set to ON.

System action: NFS client processing continues.

Operator response: Record the z/OS operator console message and notify the system programmer.

System programmer response: Check the parameters to make sure that **hard** should be ON.

GFSC315E ERROR ENCOUNTERED WHILE PARSING MOUNT PATH, REASON CODE *reasoncd***.**

Explanation: The specified mount path is not correct.

System action: The **mount** command failed.

Operator response: Record the z/OS operator console message and notify the system programmer.

System programmer response: See *z/OS UNIX System Services Messages and Codes* for a description of the reason code *reasoncd*. Correct the mount path and reissue the **mount** command.

GFSC317E ERROR ENCOUNTERED WHILE PARSING HOSTNAME, REASON CODE *reasoncd***.**

Explanation: The specified host name is not correct.

System action: The **mount** command failed.

Operator response: Record the z/OS operator console message and notify the system programmer.

System programmer response: See *z/OS UNIX System Services Messages and Codes* for a description of the reason code *reasoncd*. Correct the mount path and reissue the **mount** command.

GFSC318E READ FAILED FOR NETWORK FILE SYSTEM CLIENT MOUNT PARAMETERS.

Explanation: An error occurred while the NFS client was processing the **mount** parameters. This message follows other messages, GFSC3*xx*E, that describe the error in more detail.

System action: The **mount** command failed.

System programmer response: Correct the **mount** parameter options, and reissue the **mount** command.

GFSC319E *text* **IS AN INCORRECT KEYWORD FOR NETWORK FILE SYSTEM CLIENT INSTALLATION PARAMETER.**

Explanation: The value of *text* can be specified only as a **mount** parameter.

System action: The NFS client processing stops.

Operator response: Record the z/OS operator console message and notify the system programmer.

System programmer response: Correct the NFS client installation parameter, stop z/OS UNIX, and restart z/OS UNIX.

GFSC320E **INCORRECT OPTION** *text* **SPECIFIED FOR DELIM KEYWORD, VALID OPTION IS BINARY, CR, CRLF, CRNL, LF, LFCR, OR NL.**

Explanation: An incorrect option, *text*, has been specified for the keyword **delim**.

System action: NFS client processing stops if the error is in the NFS client installation parameter. The **mount** command failed if the error is in the **mount** parameter.

Operator response: Record the z/OS operator console message and notify the system programmer.

System programmer response: If the error is in the **mount** parameter, correct the option for the keyword **delim** and reissue the **mount** command. If the error occurs in the NFS client installation parameter, correct the option for the keyword **delim**, stop z/OS UNIX, and restart z/OS UNIX.

GFSC500I **CLIENT LOG DATA SET,** *text*, **FLUSHED.**

Explanation: The data buffer of the active client log data set, *text*, was flushed to disk.

In the message text:

text The value of *text* is the associated data set name of the active client log data set.

System action: NFS client processing continues.

GFSC501I **CLIENT LOG DATA SET** *text* **RE-INITIALIZED.**

Explanation: The error log data set is reinitialized.

In the message text:

text The value of *text* is the associated data set name of the active log data set.

System action: NFS client processing continues.

GFSC502E **CANNOT OPEN CLIENT LOG DATA SET,** *text1*, *text2*.

Explanation: The NFS client failed to open the client log data set.

In the message text:

text1 The value of *text1* is the data definition (DD) associated with the client log data set that cannot be opened.

text2 The value of *text2* is the failure reason of the C function, **fopen**.

System action: NFS client processing stops.

Operator response: Record the z/OS operator console message and notify the system programmer.

System programmer response: Fix the client log data set and reply to the z/OS UNIX message to restart the NFS client.

GFSC503E **CLIENT LOGGING ENDED.**

Explanation: The NFS client failed to manipulate the client log data set. See the previous operator console message for the failure reason.

System action: NFS client processing continues.

Operator response: Record the z/OS operator console message and notify the system programmer.

System programmer response: Fix the client log data set, stop the NFS client, and reply to the z/OS UNIX message to restart the NFS client.

GFSC504I **CLIENT LOG DATA SET SWITCHED TO** *text*.

Explanation: A `no space` or an I/O error condition was detected while the system was writing to the client log data set. NFS client logging switched to the other log data set, *text*.

System action: NFS client processing continues.

GFSC505E **MISSING DD STATEMENT OR INCORRECT DATA SET ORGANIZATION FOR LOG DATA SET.**

Explanation: The error log data set has either an incorrect data set organization or a missing data definition (DD) statement.

System action: NFS client processing stops if the error occurred during initialization. NFS client processing continues with client logging ended if the error occurred after initialization.

Operator response: Record the z/OS operator console message and notify the system programmer.

System programmer response: Correct the data

definition statement or data set organization for the error log data set, stop the NFS client, and reply to the z/OS UNIX message to restart the NFS client.

GFSC506E *text1* **FAILED FOR** *text2*, *text3*.

Explanation: The NFS client failed to manipulate the client log data set, *text2*. The value of *text3* is the failure reason for the C function *text1*.

System action: NFS client processing continues.

Operator response: Record the z/OS operator console message and notify the system programmer.

System programmer response: Fix the client log data set, stop the NFS client, and reply to the z/OS UNIX message to restart the NFS client.

GFSC700I z/OS NETWORK FILE SYSTEM CLIENT *fmid* **STARTED**

Explanation: The NFS client is initialized and ready to process NFS requests.

In the message text:

fmid The FMID of the NFS client; for example, (HDZ11VC).

System action: NFS client processing continues.

GFSC701I NETWORK FILE SYSTEM CLIENT SHUTDOWN IN PROGRESS.

Explanation: NFS client shutdown processing has started.

System action: NFS client shutdown processing continues.

GFSC702I NETWORK FILE SYSTEM CLIENT SHUTDOWN COMPLETE.

Explanation: The NFS client has completed shutdown processing.

System action: The NFS client and its associated subtasks have ended.

GFSC703E NETWORK FILE SYSTEM CLIENT INITIALIZATION FAILED: NETWORK FILE SYSTEM CLIENT IS ALREADY STARTED.

Explanation: Only one NFS client can be started on an MVS system.

System action: This NFS client initialization ends.

Operator response: Record the z/OS operator console message and notify the system programmer.

System programmer response: A prior NFS client session has not ended. UNIX end processing should have ended the NFS client colony address space. Collect installation parameters, dumps, and NFS client log data sets, and contact the IBM Support Center. Use the **force** command to end the NFS client colony address space and then restart UNIX.

GFSC704E NETWORK FILE SYSTEM CLIENT INITIALIZATION FAILED: DOWN LEVEL SECURITY PRODUCT.

Explanation: Resource Access Control Facility (RACF) is down level.

System action: NFS client processing ends.

Operator response: Record the z/OS operator console message and notify the system programmer.

System programmer response: Check to determine the required RACF level.

GFSC705E NETWORK FILE SYSTEM CLIENT INITIALIZATION FAILED: MAIN STORAGE IS UNAVAILABLE.

Explanation: The NFS client was not able to allocate the necessary storage. The cause might be that the value specified in the **bufhigh** attribute is too large or the REGION size is too small.

System action: NFS client ends.

Operator response: Record the MVS console message and notify the system programmer.

System programmer response: The requested memory is not available during NFS client initialization. Take either or both of the following actions before restarting the NFS client.

- Increase the REGION size for the client procedure.
- Decrease the value specified for the **bufhigh** attribute of the FILESYSTYPE parameter in the BPXPRM*xx* parmlib member.

GFSC707E NETWORK FILE SYSTEM CLIENT INITIALIZATION FAILED: INCORRECT PARAMETER IN INSTALLATION PARAMETERS.

Explanation: The NFS client has detected an error in the installation parameters.

System action: NFS client processing ends.

Operator response: Record the z/OS operator console message and notify the system programmer.

System programmer response: Check previous console messages prefixed with `GFSC`. Correct the parameter, stop z/OS UNIX, and restart z/OS UNIX.

GFSC708E NETWORK FILE SYSTEM CLIENT INITIALIZATION FAILED: CONVERSION SERVICE IS NOT INSTALLED OR NOT AVAILABLE. RETURN CODE *retc*, REASON CODE *rsnc*.

Explanation: Both the Unicode and character data representation architecture (CDRA) initialization requests failed.

System action: NFS client startup ends.

Operator response: Record the z/OS operator console message and notify the system programmer.

System programmer response: See *z/OS MVS System Messages, Vol 2 (ARC-ASA)*, SA22-7632 for a description of the return code *retc* and reason code *rsnc* to determine the corrective action.

GFSC709E NETWORK FILE SYSTEM CLIENT INITIALIZATION FAILED: NOT STARTED IN A STANDALONE COLONY ADDRESS SPACE.

Explanation: The NFS client detected an error during initialization processing. The NFS client was started by some means other than a UNIX kernel.

System action: NFS client processing ends.

Operator response: Record the z/OS operator console message and notify the system programmer.

System programmer response: The NFS client must be initialized by a UNIX kernel.

GFSC710E NETWORK FILE SYSTEM CLIENT INITIALIZATION FAILED: ESTAE INITIALIZATION FAILED.

Explanation: The NFS client detected an error during initialization processing.

System action: NFS client processing ends.

Operator response: Collect any dumps and NFS client log data sets, and notify the system programmer.

System programmer response: Collect installation parameters, dumps, and NFS client log data sets, and contact the IBM Support Center.

GFSC711E NETWORK FILE SYSTEM CLIENT INITIALIZATION FAILED: z/OS UNIX KERNEL SERVICE FAILED.

Explanation: The NFS client detected an error during the second phase of initialization processing.

System action: The NFS client ends.

Operator response: Collect any dumps and NFS client log data sets and notify the system programmer.

System programmer response: Collect installation parameters, dumps, and NFS client log data sets, and contact the IBM Support Center.

GFSC712E NETWORK FILE SYSTEM CLIENT INITIALIZATION FAILED: SOCKET CALL GETHOSTNAME FAILED, RETURN CODE *returncd*.

Explanation: The NFS client has detected an error (*returncd*) during initialization processing. This error might be caused by z/OS UNIX and a TCP/IP connection failure.

System action: NFS client ends.

Operator response: Collect any dumps and NFS client log data sets and notify the system programmer.

System programmer response: Collect installation parameters, dumps, and NFS client log data sets, and contact the IBM Support Center.

GFSC713E NETWORK FILE SYSTEM CLIENT LOGGING IS TERMINATED.

Explanation: The NFS client has detected an error. The client log daemon has ended.

System action: NFS client ends.

Operator response: Collect any dumps and NFS client log data sets and notify the system programmer.

System programmer response: Collect installation parameters, dumps, and NFS client log data sets, and contact the IBM Support Center.

GFSC714E NETWORK FILE SYSTEM CLIENT SVC DUMP REQUEST FAILED, RETURN CODE *returncd* REASON CODE *reasoncd*.

Explanation: A request to write an MVS SVC dump failed. See the description of the **sdump** macro in *z/OS MVS Programming: Authorized Assembler Services Reference LLA-SDU* for the meaning of return code *returncd* and reason code *reasoncd*.

System action: NFS client ends.

Operator response: Collect the MVS console log and NFS client log data sets, and notify the system programmer.

System programmer response: Collect installation parameters, dumps, and NFS client log data sets, and contact the IBM Support Center.

GFSC715E NETWORK FILE SYSTEM CLIENT ESTAE EXIT UNABLE TO COMPLETE, ABEND CODE *d_digit*.

Explanation: The NFS client recovery exit detected a recursive abend.

In the message text:

d_digit The value of *d_digit* is the last abend code encountered by a secondary instance of the ESTAE exit routine.

System action: The ESTAE processing stops. The NFS client ends.

Operator response: Collect any dumps, the z/OS console log, and NFS client log data sets, and notify the system programmer.

System programmer response: Collect installation parameters, dumps, the MVS console log, and NFS client log data sets, and contact the IBM Support Center.

GFSC716I NETWORK FILE SYSTEM CLIENT DAEMON TASK RESTARTED.

Explanation: The NFS client detected a stopped asynchronous daemon task and restarted it.

System action: The NFS client daemon has restarted. NFS client processing continues.

GFSC721E UNABLE TO SETUP ERROR RECOVERY (ESTAE), RETURN CODE *returncd***.**

Explanation: The NFS client daemon or thread failed to setup error recovery. See *z/OS MVS Programming: Authorized Assembler Services Reference ENF-IXG* for information about the **estae** macro return code *returncd*.

System action: If the error occurred in the daemon, the NFS client has initiated shutdown processing. If the error occurred in the thread, the associated operation ends and the NFS client processing continues.

Operator response: Collect any dumps, the z/OS console log, and NFS client log data sets, and notify the system programmer.

System programmer response: If the program is not in error, see the messages in the client log data sets for more information. Search problem-reporting databases for a fix for the problem. If no fix exists, contact the IBM Support Center. Provide all the printed output and copies of output data sets related to the problem.

GFSC722E A SOCKET COULD NOT BE CREATED, RETURN VALUE -1 RETURN CODE *returncd* **REASON CODE** *reasoncd***.**

Explanation: The NFS client daemon or thread processing failed to create a socket for network communications.

System action: If the error occurred in the daemon, the NFS client has initiated shutdown processing. If the error occurred in the thread, the associated operation ends and the NFS client processing continues.

Operator response: Collect any dumps, the z/OS console log, and NFS client log data sets, and notify the system programmer.

System programmer response: See *z/OS UNIX System Services Messages and Codes* for a description of return code *returncd* and reason code *reasoncd* to determine the corrective action. If the program is not in error, see the messages in the client log data sets for more information. Search problem-reporting databases for a fix for the problem. If no fix exists, contact the IBM Support Center. Provide all the printed output and copies of output data sets related to the problem.

GFSC723E NETWORK FILE SYSTEM CLIENT ABEND @ *module***+***offset***.**

Explanation: The NFS client encountered a programming error. An SVC dump was issued to capture the diagnostic information.

System action: The NFS client has initiated shutdown processing.

System programmer response: Look at the messages in the client log data sets for more information. Search problem-reporting databases for a fix for the problem. If no fix exists, contact the IBM Support Center. Provide all the printed output and copies of output data sets related to the problem

---

GFSC724E UNABLE TO BIND A RESERVED PORT TO SOCKET *socketnum***.**

Explanation: The mount daemon or umount thread failed to obtain a reserved port for socket *socketnum* for network communications.

System action: If the error occurred in the daemon, the NFS client has initiated shutdown processing. If the error occurred in the thread, the associated operation ends and the NFS client processing continues.

System programmer response: If the program is not in error, look at the messages in the client log data sets for more information. Search problem-reporting databases for a fix for the problem. If no fix exists, contact the IBM Support Center. Provide all the printed output and copies of the output data sets related to the problem.

GFSC725E SOCKET, *socketnum***, COULD NOT BE CLOSED, RETURN VALUE -1 RETURN CODE** *returncd* **REASON CODE** *reasoncd***.**

Explanation: While closing a socket *socketnum*, the system detected an error *returncd* was detected.

System action: The NFS client continues processing.

System programmer response: See *z/OS UNIX System Services Messages and Codes* for a description of return code *returncd* and reason code *reasoncd* to determine the corrective action. If the program is not in error, look at

the messages in the client log data sets for more information. Search problem-reporting databases for a correction for the error. If no fix exists, contact the IBM Support Center. Provide all the printed output and copies of output data sets related to the problem.

GFSC726E REQUEST (*requestid*) : THE PORTMAPPER OF THE SERVER *hostname* DOES NOT RESPOND TO PMAP_GETMAPS

Explanation: The **mount** request with code *requestid* failed because the user specified **proto(tcp)** and the *hostname* did not publish its registered remote program port for the remote procedure call (RPC).

System action: The **mount** command ended with an error. NFS client continues processing.

User response: Verify that the server does not respond to **orpcinfo**. Reissue the **mount** command without **proto(tcp)**.

GFSC727W REQUEST (*requestid*) : THE PORTMAPPER OF THE SERVER *hostname* DOES NOT RESPOND TO PMAP_GETMAPS. ATTEMPTING UDP RPC WITH NFS VERSION *d_digit* PROTOCOL.

Explanation: The NFS client warns the user who mounts to *hostname* that the host might not have its portmapper running; the NFS client uses the user datagram protocol (UDP) remote procedure call (RPC). A UDP RPC is attempted with the NFS protocol.

System action: The **mount** command might or might not succeed. NFS client continues processing.

User response: Verify that the server does not respond to the code **orpcinfo**. If the **mount** command subsequently fails, then the host likely does not have the NFS server.

GFSC728E REQUEST (requestid) : UNABLE TO CONNECT SOCKET=*socketno* PORT=*portno1* WITH THE SERVER hostname PORT=*portno2*

Explanation: The NFS client uses TCP/IP to connect socket *socketno* with port *portno1* to the server *hostname* with port *portno2*; however, the connect system call failed.

System action: The operation ended with an error. The NFS client processing continues.

User response: Verify that the NFS server on *hostname* is available and running. Verify that the TCP/IP subsystem on z/OS is available and running. If the **mount** command subsequently fails, then the host likely does not have the NFS server.

GFSC729W z/OS NFS CLIENT IS UNABLE TO RESERVE *d_digit* SOCKETS, CURRENT MAXFILEPROC IS *d_digit*

Explanation: The maximum number of socket descriptors that a process can have open concurrently has been exceeded. Increase the MAXFILEPROC parameter value in the BPXPRM*xx* member to bypass this problem. See *z/OS UNIX System Services Planning* for more information about specifying the MAXFILEPROC value.

System action: The operation ended with an error. The NFS client processing continues.

System programmer response: Increase the MAXFILEPROC parameter value in the BPXPRM*xx* member to bypass the problem.

GFSC840I usage: *text* [-a] [-d] [-e] [host]

Explanation: This is the usage for the **showmount** command. The value of *text* is the command as entered by the user. The valid options follow:

-a Display all mounts in the format *hostname:directory* from host NFS server

-d Display only directory names of all mounts from host NFS server

-e Display the list of exported directories from host NFS server

---

GFSC841E Unknown host *text*

Explanation: The user entered incorrect host address information, *text*.

System action: The system stops processing the command.

User response: Correct the syntax and reissue the command.

GFSC842E Cannot resolve local host name

Explanation: The local host name was not found.

System action: The system stops processing the command.

User response: Contact the system administrator to check the TCP/IP configuration.

GFSC843E Unknown flag '-*character*'

Explanation: An incorrect option, '*-character*', was specified.

System action: The system stops processing the command.

User response: Correct the syntax and reissue the command.

GFSC845I **usage:** *text input output*

Explanation: This is the usage for the **crnl2nl** and **nl2crnl** commands.

In the message text:

text The value of *text* is the command as entered by the user.

The valid parameters follow:

input Absolute path name of the input file to be converted.

output Absolute path name of the output file.

GFSC846E **Cannot open input file,** *text1*:*text2*

Explanation: The system cannot open the input file, *text1*.

In the message text:

text1 The value of *text1* is the input path name entered by the user.

text2 The value of *text2* is the failure information returned during an attempt to open the input file.

System action: The system stops processing the command.

User response: Check the input file, *text1*.

GFSC847E **Cannot open output file,** *text1*:*text2*

Explanation: The system cannot open output file, *text1*.

In the message text:

text1 The value of *text1* is the output path name entered by the user.

text2 The value of *text2* is the failure information returned during an attempt to open the output file.

System action: The system stops processing the command.

User response: Check the output file, *text1*.

GFSC848E **Cannot read input file,** *text1*: *text2*

Explanation: The system cannot read the input file, *text1*.

In the message text:

text1 The value of *text1* is the input path name entered by the user.

text2 The value of *text2* is the failure information returned during an attempt to read the input file.

System action: The system stops processing the command.

User response: Check the input file, *text1*.

GFSC849E **Cannot write output file,** *text1*: *text2*

Explanation: The system cannot write to the output file, *text1*.

In the message text:

text1 The value of *text1* is the output path name entered by the user.

text2 The value of *text2* is the failure information returned during an attempt to write to the output file.

System action: The system stops processing the command.

User response: Check the output file, *text1*.

GFSC850E **Input path name cannot be equal to output path name.**

Explanation: The input path name cannot be the same as the output path name.

System action: The system stops processing the command.

User response: Correct the syntax and reissue the command.

GFSC854I **usage:** *text* **[-crnzm <mount point>]**

Explanation: This is the usage for the **nfsstat** command.

In the message text:

text The value of *text* is the command entered by the user.

The valid parameters follow:

-c Display both NFS and remote procedure call (RPC) statistics about the NFS client.

-n Display NFS statistics about the NFS client.

-r Display remote procedure call (RPC) statistics about the NFS client.

-z Initializes statistics to zero. This option is for use by the root user only and can be combined with any of the preceding options. Zero a particular set of statistics after printing them.

-m Display the name of each NFS-mounted file system.

-m mount point
Display information about the NFS mounted file system on the specified mount point.

GFSC855E Must be a root user to issue '*character*' flag

Explanation: The option '*character*' can be issued only with the root authority.

System action: The system stops processing the command.

User response: Contact your system administrator to issue this command.

GFSC856E Network File System Client command, *text*, failed, return value -1 return code *returncd* reason code *reasoncd*

Explanation: The command *text* failed.

System action: The system stops processing the command.

User response: See *z/OS UNIX System Services Messages and Codes* for a description of the return code *returncd*. See Table 67 on page 309 for more information about the reason code *reasoncd*.

GFSC858E Directory *text* not mounted.

Explanation: The directory *text* was not mounted.

System action: The system stops processing the command.

User response: Issue **nfsstat -m** to view the list of active mount points. If the mount point does not exist, contact the system administrator to mount the directory.

Messages from the client platform (AIX)

This section provides a list of messages from the client operating system, in response to NFS reply results (messages without message numbers). These messages are platform dependent.

Note: The message text is from an AIX system.

Cross device link

Explanation:

1. An attempt has been made to rename a member of a partitioned data set (PDS) or partitioned data set extended (PDSE), but the target file is not a member of the same PDS or PDSE.
2. An attempt has been made to rename a non-PDS or non-PDSE file, but the target file is a member of a PDS or PDSE.

User response: Try a copy and remove operation, instead of a rename operation.

Directory Not Empty

Explanation: An attempt was made to remove a PDS or PDSE with members.

User response: Delete all members before trying to remove a PDS or PDSE.

File exists

Explanation: An attempt was made to rename a PDS or PDSE, but the target file already exists.

User response: Delete the target file before renaming the PDS or PDSE. This is not required for a regular file or for a PDS or PDSE member.

File Name Too Long

Explanation: The file name is not a valid z/OS data set or member name.

User response: File names must follow the z/OS naming conventions. See *z/OS DFSMS Using Data Sets* for data set naming conventions.

Invalid

Explanation:

1. The specified parameters were incorrect.
2. The creation of a virtual storage access method (VSAM) data set failed.

User response: Respecify the correct parameters, or contact the system programmer to determine the VSAM data set failure on the NFS server.

I/O Error (with possible system programmer response)

Explanation:

1. An unexpected error from catalog management occurred.
2. Dynamic allocation failed during an action other than a read or write.
3. A file could not be opened during an action other than a read or write.
4. An error occurred during the reading of a PDS or PDSE directory.
5. No space is available in the task input/output table (TIOT).
6. The TIOT resource is unavailable.
7. The system is unable to release enough resources.
8. Insufficient units are available.
9. The server could not get enough memory to perform this function.

System programmer response: For explanation 9, stop the server and change the region field in the job control language (JCL) before restarting, or modify parameters in the attributes file. For the other explanations, perform the appropriate action.

I/O Error (with possible user response)

Explanation:

1. An attempt was made to nest PDSs or PDSEs.
2. The maximum number of file allocations was exceeded.
3. The file is being written in text mode, the new write request offset is determined to fall within the end-of-line (EOL) sequence (**lf**, **cr**, **crlf**, **lfcr**) of a previous line, and the new data does not contain the correct EOL characters.
4. The file is being written in text mode, with a nonzero EOL (**lf**, **cr**, **lfcr**, **crlf**). The number of bytes of data in the written line is larger than the maximum record size of the file.
5. The file is being written in text mode, with fixed records, a nonzero EOL, and **blankstrip** not set (no padding blanks on write), and the number of bytes of data in the written line is less than or greater than the record size of the file.
6. The file is being written in text mode, with fixed records and **blankstrip** set, and the line of written data contains trailing blanks as part of the data.
7. When a file containing a 0 length line is written to z/OS as **recfm(u)** in text mode, a write error occurs.

8. An access method services alias name was specified in a remove (**rm** or **rmdir**) or rename (**mv**) request.
9. An "s" was specified in the **recfm** attribute for a PDS or PDSE.
10. If you try to append data to a member of a PDS or PDSE, an I/O error occurs.
11. An incorrect attribute was specified in the command.

User response: Perform the appropriate action.

Is a directory

Explanation: A nondirectory operation has been tried on a PDS or PDSE.

User response: Use directory operations on the data set.

NFS server *name* **not responding still trying**

Explanation: Long delays occurred between operations.

User response: Possible user responses follow:

- The server might need extra time to service client requests. Wait until the message "`NFS server` *name* `ok`" appears.
- Determine if the network traffic is heavy and overloaded. Try to isolate the path where the client workstation communicates with the server machine.
- Determine if the client workstation transmits the same requests over and over. Commands such as **nfsstat -c** on AIX or UNIX platforms show the number of client retransmissions (**retrans**) as well as the number of **badcalls** and **badxid**. When the values of **badcalls** and **badxid** are high, the client machine usually has bad retransmissions. High retransmissions might be caused by an overloaded network or a slow server.
- If the network is overloaded, contact the network administrator.
- If the server is slow, determine if the client workstation tends to transmit requests out of sequence or incompletely, and you are performing I/O to a file in text mode. If the request is incomplete, the client might not send the remainder of the request until a later time. In this case, increase the value of the server's **cachewindow** attribute. The **cachewindow** buffers store out of sequence and incomplete requests from clients. The general rule in setting the *n* value of cachewindow(*n*) is

  ```
  n  = (( num of BIOD  + 1 ) *
     (client_max_IO_buffer_size/transfer_size))
  ```

  In the formula:

  num of BIOD
  Is the number of blocked I/O daemons set by the client workstation. This value is usually set to defaults at the installation of the operating system or by the system administrator.

  client_max_IO_buffer_size
  Is the amount of I/O data that is requested by the client (for example, the client writes 8192 bytes of data to the remote file system). This value is determined by the application programs.

  transfer_size
  Is the actual size of data being sent across the network (for example, the 8192 bytes of data can be broken down into 16 smaller packets of 512 bytes (16x512=8192)). This value is determined dynamically by the client workstation.
- If your client workstation tends to send duplicate transmissions for the same request too often, thus increasing the workload on the server, you might want to decrease the client's retransmission rate and increase the request timeout. In the **mount** command, you can specify these values:

  ```
  mount -o retrans=3,timeout=30 IO_buffer_size
  ```

  In the command:

  retrans
  Is the number of retransmissions allowed before a timeout. The default value is vendor specific, ranging from 3 to 5.

  timeout
  Is the timeout value in tenths of a second. timeout=30 means 3 seconds. The default value is vendor specific, ranging from 5 to 11.

  If the reply is not received by the client within the timeout period, a minor timeout occurred for this request. The timeout period is doubled, and the request is resent. The process is repeated until the retransmission count specified by **retrans** is reached; if no reply is received, a major timeout occurred.

No such device

Explanation:

Step 1. The file resides on an off-line device

Step 2. The file was migrated to another storage level. Whether the file is being recalled depends on the **retrieve** attribute

User response: For Step 1, contact the system operator or system programmer.

For Step 2, try the request again later, if the **retrieve** attribute is enabled. If not, try the request again, with the **retrieve** attribute enabled.

No space left on device

Explanation:

1. The file exceeded the space that was allocated for it.
2. The PDS or PDSE exceeded the space that was allocated for it.
3. The PDS directory exceeded the space that was allocated for it.

User response:

1. For explanation 1, save the file into a larger file, and then rename it to the old name, if necessary.
2. For explanation 2, create a larger PDS or PDSE and store the member there. Then copy the members of the old PDS or PDSE to the new PDS or PDSE. Rename the new PDS or PDSE, if necessary.
3. For explanation 3, create a new PDS with a larger directory (use the **dir** attribute). Store the member in the new PDS. Then copy the members from the old PDS to the new PDS and rename them, if necessary.

No such file or directory

Explanation: A **locate** command failed for this file. The file is not cataloged, or the system operator might have unmounted the file before you issued the **umount** command.

User response: Check the spelling. If the spelling is correct, contact the system administrator.

Not a directory

Explanation: A directory operation was made on a file that is not a PDS or PDSE.

User response: Use nondirectory operations on the file.

Not Owner

Explanation:

1. File has not timed out yet.
2. File is open. Another client (maybe even the same client) has the file open for writing.
3. The client tried to change the mode in a **nfsattr** NFS procedure call.

User response: Possible user responses follow:

- For explanation 1, follow the steps in waiting and retrying that are described under the "`Permission denied`" message,
- For explanation 3, do not try to change the mode.

Permission denied

Explanation:

1. The file is not in use but has not timed out yet (occurs most often when writing to PDS or PDSE members without **writetimeout** expiring between saving members).
2. The file is in use by a z/OS user or another client.
3. Dynamic allocation–authorized function requested by an unauthorized user. The error codes from dynamic allocation are printed to the log data set.
4. The Resource Access Control Facility (RACF) is not active.
5. Dynamic allocation–The request was denied by the operator (dsname allocation). The error codes from dynamic allocation are printed to the log data set.
6. Dynamic allocation–The installation validation routine denied this request. The error codes from dynamic allocation are printed to the log data set.
7. You are not authorized for this request.
8. The system operator suspended mount request processing (only if this message appears following a mount attempt).
9. The file/prefix is not exported to the client (only if this message appears following a mount attempt).
10. IDCAMS failed during a rename or remove procedure. Usually this happens because the file is in use. The output from IDCAMS is printed to the log data set.
11. With OS/2®, you might get a `"SYS0055 Access denied"` message if the **noretrieve** attribute is set and a **dir** command is done against a mounted file system containing migrated files.

User response: For explanation 1, to determine if this is the problem, check the timeout values (**attrtimeout**, **readtimeout**, **writetimeout**). Retry the request after the shortest timeout expires. If the request still fails, retry the request after the next shortest timeout expires. If the request still fails, retry after the longest timeout expires. If the request still fails, this is not the problem.

For explanations 2 and 10, retry the request later.

For explanations 3, 4, 5, 6, and 9, notify the system programmer.

For explanation 7, enter the **mvslogin** command again and retry the request. If the request still fails, notify your system programmer.

For explanation 8, notify your system operator or system programmer.

For explanation 11, specify the **retrieve** attribute in the **mount** command, or the system administrator can make that the default.

Read Only File System

Explanation: One of these NFS procedures was tried on a read-only file system: **link**, **write**, **rename**, **remove**, **mkdir**, or **create**.

User response: See the documentation on the exports data set to see how a file system is designated read only (see *z/OS Network File System Guide and Reference*). The exports data set needs to be changed, or you are using it incorrectly.

Stale NFS File Handle

Explanation: A file handle is used by the client and server sides of the NFS to specify a particular file or prefix. A stale file handle occurs when the name is no longer valid, possibly due to one of the these conditions:

1. The file or prefix was removed by the system operator.
2. The server was stopped and restarted. This affects files and members below mount points.

User response: For explanation 1, unmount and mount again. If the client maintains that the device is busy even though it is not, you might have to restart the client.

For explanation 2, enter the **mvslogin** command again and retry the request.

Weak Authorization

Explanation: The authorization data in the remote procedure call (RPC) message was not valid. This is a client side error.

User response: UNIX style authorization is required.

---

value to large

Explanation: The NFS version 4 protocol defines 64-bit file ids which may be returned to the client application as a 64-bit value. Client applications may receive err79 ″Value too large″ when using 32-bit system calls such as UNIX stat() against remote NFS version 4 server objects. Note: this does not apply to z/OS NFS client applications because NFS version 4 is not supported in the z/OS V1R7 client.

User response: Change your application to use 64-bit system calls such as UNIX stat64() to prevent err79 messages in client applications when using NFS version 4.

Chapter 20. Return codes

Table 61 lists the externalized return codes that are defined by the NFS version 2 protocol.

Table 61. Externalized return codes defined by the NFS version 2 protocol

Return Value	Return Code	Description
NFS_OK	0	Requests completed successfully and the results are valid.
NFSERR_PERM	1	Not owner. The caller does not have correct ownership to perform the requested operation.
NFSERR_NOENT	2	No such file or directory. The file or directory specified does not exist.
NFSERR_IO	5	A hard error occurred when the operation was in progress. For example, this could be a disk error.
NFSERR_NXIO	6	No such device or address.
NFSERR_ACCESS	13	Permission denied. The caller does not have the correct permission to perform the requested operation.
NFSERR_EXIST	17	File exists. The file specified already exists.
NFSERR_NODEV	19	No such device.
NFSERR_NOTDIR	20	Not a directory. The caller specified a non-directory in a directory operation.
NFSERR_ISDIR	21	Is a directory. The caller specified a directory in a non-directory operation.
NFSERR_EINVAL	22	An argument was passed to the z/OS NFS server was not valid.
NFSERR_FBIG	27	File too large. The operation caused a file to grow beyond the server's limit.
NFSERR_NOSPC	28	No space left on device. The operation caused the server's file system to reach its limit.
NFSERR_ROFS	30	Read-only file system. Write tried on a read-only file system.
NFSERR_NAMETOOLONG	63	File name too long. The file name in an operation was too long.
NFSERR_NOTEMPTY	66	Directory not empty. Tried to remove a directory that was not empty.
NFSERR_DQUOT	69	Disk quota exceeded. The client's disk quota on the server has been exceeded.
NFSERR_STALE	70	The file handle given in the arguments was not valid. That is, the file referred to by that file handle no longer exists, or access to it has been revoked.

Table 62 lists the externalized return codes that are defined by the NFS version 3 protocol.

Table 62. Externalized return codes defined by the NFS version 3 protocol

Return Value	Return Code	Description
NFS_OK	0	Requests completed successfully and the results are valid.
NFS3ERR_PERM	1	Not owner. The caller does not have correct ownership to perform the requested operation.
NFS3ERR_NOENT	2	No such file or directory. The file or directory specified does not exist.
NFS3ERR_IO	5	A hard error occurred when the operation was in progress. For example, this could be a disk error.
NFS3ERR_NXIO	6	No such device or address.

© Copyright IBM Corp. 1991, 2006

Table 62. Externalized return codes defined by the NFS version 3 protocol (continued)

Return Value	Return Code	Description
NFS3ERR_ACCESS	13	Permission denied. The caller does not have the correct permission to perform the requested operation.
NFS3ERR_EXIST	17	File exists. The file specified already exists.
NFS3ERR_XDEV	18	Attempt to do an operation across the file system.
NFS3ERR_NODEV	19	No such device.
NFS3ERR_NOTDIR	20	Not a directory. The caller specified a non-directory in a directory operation.
NFS3ERR_ISDIR	21	Is a directory. The caller specified a directory in a non-directory operation.
NFS3ERR_INVAL	22	An argument was passed to the z/OS NFS server was not valid.
NFS3ERR_FBIG	27	File too large. The operation caused a file to grow beyond the server's limit.
NFS3ERR_NOSPC	28	No space left on device. The operation caused the server's file system to reach its limit.
NFS3ERR_ROFS	30	Read-only file system. Write tried on a read-only file system.
NFS3ERR_MLINK	31	Too many links.
NFS3ERR_NAMETOOLONG	63	File name too long. The file name in an operation was too long.
NFS3ERR_NOTEMPTY	66	Directory not empty. Tried to remove a directory that was not empty.
NFS3ERR_DQUOT	69	Disk quota exceeded. The client's disk quota on the server has been exceeded.
NFS3ERR_STALE	70	The file handle given in the arguments was not valid. That is, the file referred to by that file handle no longer exists, or access to it has been revoked.
NFS3ERR_NOT_SYNC	10001	File handle is not valid.
NFS3ERR_BAD_COOKIE	10002	Synchronization mismatch on SETATTR.
NFS3ERR_BAD_COOKIE	10003	READDIR and READDIRPLUSS cookie is stale.
NFS3ERR_NOTSUPP	10004	Operation is not supported.
NFS3ERR_TOOSMALL	10005	Buffer or request is too small.
NFS3ERR_SERVERFAULT	10006	Server abandons the request.
NFS3ERR_BADTYPE	10007	Type of an object is not supported.
NFS3ERR_JUKEBOX	10008	Request was initiated, but not completed.

Table 63 lists the externalized return codes that are defined by the NFS version 4 protocol.

Table 63. Externalized return codes defined by the NFS version 4 protocol

Return Value	Return Code	Description
NFS4_OK	0	Requests completed successfully and the results are valid.
NFS4ERR_PERM	1	Not owner. The caller does not have correct ownership to perform the requested operation.
NFS4ERR_NOENT	2	No such file or directory. The file or directory specified does not exist.
NFS4ERR_IO	5	A hard error occurred when the operation was in progress. For example, this could be a disk error.
NFS4ERR_NXIO	6	No such device or address.
NFS4ERR_ACCESS	13	Permission denied. The caller does not have the correct permission to perform the requested operation.
NFS4ERR_EXIST	17	File exists. The file specified already exists.

Table 63. Externalized return codes defined by the NFS version 4 protocol (continued)

Return Value	Return Code	Description
NFS4ERR_XDEV	18	Attempt to do an operation across the file system.
	19	Not used/reserved.
NFS4ERR_NOTDIR	20	Not a directory. The caller specified a non-directory in a directory operation.
NFS4ERR_ISDIR	21	Is a directory. The caller specified a directory in a non-directory operation.
NFS4ERR_INVAL	22	An argument was passed to the z/OS NFS server was not valid.
NFS4ERR_FBIG	27	File too large. The operation caused a file to grow beyond the server's limit.
NFS4ERR_NOSPC	28	No space left on device. The operation caused the server's file system to reach its limit.
NFS4ERR_ROFS	30	Read-only file system. Write tried on a read-only file system.
NFS4ERR_MLINK	31	Too many hard links.
NFS4ERR_NAMETOOLONG	63	File name too long. The file name in an operation was too long.
NFS4ERR_NOTEMPTY	66	Directory not empty. Tried to remove a directory that was not empty.
NFS4ERR_DQUOT	69	Disk quota exceeded. The client's disk quota on the server has been exceeded.
NFS4ERR_STALE	70	The file handle given in the arguments was not valid. That is, the file referred to by that file handle no longer exists, or access to it has been revoked.
NFS4ERR_BADHANDLE	10001	File handle is not valid.
NFS4ERR_BAD_COOKIE	10003	READDIR and READDIRPLUSS cookie is stale.
NFS4ERR_NOTSUPP	10004	Operation is not supported.
NFS4ERR_TOOSMALL	10005	Buffer or request is too small.
NFS4ERR_SERVERFAULT	10006	Server abandons the request.
NFS4ERR_BADTYPE	10007	Type of an object is not supported.
NFS4ERR_DELAY	10008	Request was initiated, but not completed. File was busy; retry.
NFS4ERR_SAME	10009	Request was initiated, but attributes are the same.
NFS4ERR_DENIED	10010	Lock was unavailable.
NFS4ERR_EXPIRED	10011	Lock release expired.
NFS4ERR_LOCKED	10012	I/O failure due to lock.
NFS4ERR_GRACE	10013	In grace period.
NFS4ERR_FHEXPIRED	10014	File handle expired.
NFS4ERR_SHARE_DENIED	10015	Share reserve denied.
NFS4ERR_WRONGSEC	10016	Wrong security level.
NFS4ERR_CLID_INUSE	10017	Client id in use.
NFS4ERR_RESOURCE	10018	Resource exhaustion.
NFS4ERR_MOVED	10019	Filesystem relocated.
NFS4ERR_NOFILEHANDLE	10020	Current file handle is not set.
NFS4ERR_MINOR_VERS_MISMATCH	10021	Minor version not supported.
NFS4ERR_STALE_CLIENTID	10022	Server has rebooted.
NFS4ERR_STALE_STATEID	10023	Server has rebooted.
NFS4ERR_OLD_STATEID	10024	State is not in synch.
NFS4ERR_BAD_STATEID	10025	Incorrect state id.
NFS4ERR_BAD_SEQID	10026	Request is out of sequence.
NFS4ERR_NOT_SAME	10027	Verify - attributes not the same.

Table 63. Externalized return codes defined by the NFS version 4 protocol (continued)

Return Value	Return Code	Description
NFS4ERR_LOCK_RANGE	10028	Lock range not supported.
NFS4ERR_SYMLINK	10029	Should be file/directory.
NFS4ERR_RESTOREFH	10030	No saved file handle.
NFS4ERR_LEASE_MOVED	10031	Some file system moved.
NFS4ERR_ATTRNOTSUPP	10032	Recommended attribute not supported.
NFS4ERR_NO_GRACE	10033	Reclaim attempt was not within the grace period.
NFS4ERR_RECLAIM_BAD	10034	Reclaim error occurred at server.
NFS4ERR_RECLAIM_CONFLICT	10035	Conflict occurred on reclaim.
NFS4ERR_BADXDR	10036	XDR decode failed.
NFS4ERR_LOCKS_HELD	10037	File locks held at CLOSE.
NFS4ERR_OPENMODE	10038	Conflict in OPEN and I/O.
NFS4ERR_BAD_OWNER	10039	Owner translation was not correct.
NFS4ERR_BADCHAR	10040	utf-8 character was not supported.
NFS4ERR_BADNAME	10041	Name is not supported.
NFS4ERR_BAD_RANGE	10042	Lock range was not supported.
NFS4ERR_LOCK_NOTSUPP	10043	No atomic upgrade or downgrade.
NFS4ERR_OP_ILLEGAL	10044	Operation was undefined.
NFS4ERR_DEADLOCK	10045	Deadlock occurred in file locking.
NFS4ERR_FILE_OPEN	10046	An open file blocked the operation.
NFS4ERR_ADMIN_REVOKED	10047	The lock owner's state was revoked.
NFS4ERR_CB_PATH_DOWN	10048	The callback path was down.

Table 64 lists the z/OS UNIX return codes and their equivalent NFS Version 2 server return codes.

Table 64. z/OS NFS Server: z/OS UNIX return codes mapped to NFS Version 2 return codes

OMVS Codes	Dec	Hex	NFS V2 Codes (NFSERR_)	Dec	Hex	Description
EACCES	111	006F	ACCES	13	000D	Permission denied
EAGAIN	112	0070	IO	5	0005	Resource is temporarily unavailable
EBUSY	114	0072	IO	5	0005	Resource is busy
EDEADLK	116	0074	IO	5	0005	A resource deadlock is avoided
EEXIST	117	0075	EXIST	17	0011	The file exists
EFAULT	118	0076	IO	5	0005	The address is incorrect
EFBIG	119	0077	FBIG	27	001B	The file is too large
EINVAL	121	0079	IO	5	0005	The parameter is incorrect
EIO	122	007A	IO	5	0005	An I/O error occurred

Table 64. z/OS NFS Server: z/OS UNIX return codes mapped to NFS Version 2 return codes (continued)

OMVS Codes	Dec	Hex	NFS V2 Codes (NFSERR_)	Dec	Hex	Description
EISDIR	123	007B	ISDIR	21	0015	The file specified is a directory
EMFILE	124	007C	IO	5	0005	Too many files are open for this directory
EMLINK	125	007D	IO	5	0005	Too many links occurred
ENAMETOOLONG	126	007E	NAMETOOLONG	63	003F	The file name is too long
ENFILE	127	007F	IO	5	0005	Too many files are open
ENODEV	128	0080	NODEV	19	0013	No such device exists
ENOENT	129	0081	NOENT	2	0002	No such file, directory or IPC member exists
ENOMEM	132	0084	IO	5	0005	Not enough space is available
ENOSPC	133	0085	NOSPC	28	001C	No space is left on device
ENOTDIR	135	0087	NOTDIR	20	0014	Not a directory
ENOTEMPTY	136	0088	NOTEMPTY	66	0042	Directory is not empty
ENXIO	138	008A	NXIO	6	0006	No such device or address exists
EPERM	139	008B	IO	5	0005	The operation is not permitted
EROFS	141	008D	ROFS	30	001E	The specified file system is read only
EXDEV	144	0090	XDEV	18	0012	A link to a file on another file system was attempted
E2BIG	145	0091	IO	5	0005	The parameter list is too long
ELOOP	146	0092	IO	5	0005	A loop is encountered in symbolic links
EILSEQ	147	0093	IO	5	0005	The byte sequence is illegal
EMVSERR	157	009D	IO	5	0005	MVS environmental or internal error
EMVSPARM	158	009E	IO	5	0005	Bad parameters were passed to the service

Table 64. z/OS NFS Server: z/OS UNIX return codes mapped to NFS Version 2 return codes (continued)

OMVS Codes	Dec	Hex	NFS V2 Codes (NFSERR_)	Dec	Hex	Description
EMVSPFSFILE	159	009F	IO	5	0005	z/OS UNIX encountered a permanent file error
EMVSPFSPERM	162	00A2	IO	5	0005	z/OS UNIX encountered a system error
EMVSSAFEXTRERR	163	00A3	IO	5	0005	SAF/RACF extract error
EMVSSAF2ERR	164	00A4	IO	5	0005	SAF/RACF error
EDQUOT	1133	046D	DQUOT	69	0045	Disk quota exceeded
ESTALE	1134	046E	STALE	70	0046	Stale NFS file handle
EREMOTE	1135	046F	IO	5	0005	Too many levels of remote in path

Table 65 lists the z/OS UNIX return codes and their equivalent NFS Version 3 server return codes.

Table 65. z/OS NFS Server: z/OS UNIX return codes mapped to NFS Version 3 return codes

OMVS Codes	Dec	Hex	NFS V3 Codes (NFS3ERR_)	Dec	Hex	Description
EACCES	111	006F	ACCES	13	000D	Permission denied
EAGAIN	112	0070	JUKEBOX	10008	2718	Resource is temporarily unavailable
EBUSY	114	0072	JUKEBOX	10008	2718	Resource is busy
EDEADLK	116	0074	DEADLOCK	10045	273D	A resource deadlock is avoided
EEXIST	117	0075	EXIST	17	0011	The file exists
EFAULT	118	0076	SERVERFAULT	10006	2716	The address is incorrect
EFBIG	119	0077	FBIG	27	001B	The file is too large
EINVAL	121	0079	INVAL	22	0016	The parameter is incorrect
EIO	122	007A	IO	5	0005	An I/O error occurred
EISDIR	123	007B	ISDIR	21	0015	The file specified is a directory
EMFILE	124	007C	JUKEBOX	10008	2718	Too many files are open for this directory

Table 65. z/OS NFS Server: z/OS UNIX return codes mapped to NFS Version 3 return codes (continued)

OMVS Codes	Dec	Hex	NFS V3 Codes (NFS3ERR_)	Dec	Hex	Description
EMLINK	125	007D	MLINK	31	001F	Too many links occurred
ENAMETOOLONG	126	007E	NAMETOOLONG	63	003F	The file name is too long
ENFILE	127	007F	DQUOT	69	0045	Too many files are open
ENODEV	128	0080	NODEV	19	0013	No such device exists
ENOENT	129	0081	NOENT	2	0002	No such file, directory or IPC member exists
ENOMEM	132	0084	JUKEBOX	10008	2718	Not enough space is available
ENOSPC	133	0085	NOSPC	28	001C	No space is left on device
ENOTDIR	135	0087	NOTDIR	20	0014	Not a directory
ENOTEMPTY	136	0088	NOTEMPTY	66	0042	Directory is not empty
ENXIO	138	008A	NXIO	6	0006	No such device or address exists
EPERM	139	008B	PERM	1	0001	The operation is not permitted
EROFS	141	008D	ROFS	30	001E	The specified file system is read only
EXDEV	144	0090	XDEV	18	0012	A link to a file on another file system was attempted
E2BIG	145	0091	SERVERFAULT	10006	2716	The parameter list is too long
ELOOP	146	0092	IO	5	0005	A loop is encountered in symbolic links
EMVSERR	157	009D	SERVERFAULT	10006	2716	MVS environmental or internal error
EMVSPARM	158	009E	SERVERFAULT	10006	2716	Bad parameters were passed to the service
EMVSPFSFILE	159	009F	IO	5	0005	z/OS UNIX encountered a permanent file error
EMVSPFSPERM	162	00A2	IO	5	0005	z/OS UNIX encountered a system error

Table 65. z/OS NFS Server: z/OS UNIX return codes mapped to NFS Version 3 return codes (continued)

OMVS Codes	Dec	Hex	NFS V3 Codes (NFS3ERR_)	Dec	Hex	Description
EMVSSAFEXTRERR	163	00A3	IO	5	0005	SAF/RACF extract error
EMVSSAF2ERR	164	00A4	IO	5	0005	SAF/RACF error
EDQUOT	1133	046D	DQUOT	69	0045	Disk quota exceeded
ESTALE	1134	046E	STALE	70	0046	Stale NFS file handle
EREMOTE	1135	046F	REMOTE	71	0077	Too many levels of remote in path

Table 66 lists the z/OS UNIX return codes and their equivalent NFS Version 4 server return codes.

Table 66. z/OS NFS Server: z/OS UNIX return codes mapped to NFS Version 4 return codes

OMVS Codes	Dec	Hex	NFS V4 Codes (NFS4ERR_)	Dec	Hex	Description
EACCES	111	006F	ACCES	13	000D	Permission denied
EAGAIN	112	0070	DELAY	10008	2718	Resource is temporarily unavailable
EBUSY	114	0072	DELAY	10008	2718	Resource is busy
EDEADLK	116	0074	DEADLOCK	10045	273D	A resource deadlock is avoided
EEXIST	117	0075	EXIST	17	0011	The file exists
EFAULT	118	0076	SERVERFAULT	10006	2716	The address is incorrect
EFBIG	119	0077	FBIG	27	001B	The file is too large
EINVAL	121	0079	INVAL	22	0016	The parameter is incorrect
EIO	122	007A	IO	5	0005	An I/O error occurred
EISDIR	123	007B	ISDIR	21	0015	The file specified is a directory
EMFILE	124	007C	DQUOT	69	0045	Too many files are open for this directory
EMLINK	125	007D	MLINK	31	001F	Too many links occurred
ENAMETOOLONG	126	007E	NAMETOOLONG	63	003F	The file name is too long
ENFILE	127	007F	DQUOT	69	0045	Too many files are open
ENODEV	128	0080	NXIO	6	0006	No such device exists

Table 66. z/OS NFS Server: z/OS UNIX return codes mapped to NFS Version 4 return codes (continued)

OMVS Codes	Dec	Hex	NFS V4 Codes (NFS4ERR_)	Dec	Hex	Description
ENOENT	129	0081	NOENT	2	0002	No such file, directory or IPC member exists
ENOMEM	132	0084	DELAY	10008	2718	Not enough space is available
ENOSPC	133	0085	NOSPC	28	001C	No space is left on device
ENOTDIR	135	0087	NOTDIR	20	0014	Not a directory
ENOTEMPTY	136	0088	NOTEMPTY	66	0042	Directory is not empty
ENXIO	138	008A	NXIO	6	0006	No such device or address exists
EPERM	139	008B	PERM	1	0001	The operation is not permitted
EROFS	141	008D	ROFS	30	001E	The specified file system is read only
EXDEV	144	0090	XDEV	18	0012	A link to a file on another file system was attempted
E2BIG	145	0091	SERVERFAULT	10006	2716	The parameter list is too long
ELOOP	146	0092	IO	5	0005	A loop is encountered in symbolic links
EMVSERR	157	009D	SERVERFAULT	10006	2716	MVS environmental or internal error
EMVSPARM	158	009E	SERVERFAULT	10006	2716	Bad parameters were passed to the service
EMVSPFSFILE	159	009F	IO	5	0005	z/OS UNIX encountered a permanent file error
EMVSPFSPERM	162	00A2	IO	5	0005	z/OS UNIX encountered a system error
EMVSSAFEXTRERR	163	00A3	IO	5	0005	SAF/RACF extract error
EMVSSAF2ERR	164	00A4	IO	5	0005	SAF/RACF error
EMVSNORTL	167	00A7	IO	5	0005	Access to z/OS UNIX version of C RTL denied
EDQUOT	1133	046D	DQUOT	69	0045	Disk quota exceeded

Table 66. z/OS NFS Server: z/OS UNIX return codes mapped to NFS Version 4 return codes (continued)

OMVS Codes	Dec	Hex	NFS V4 Codes (NFS4ERR_)	Dec	Hex	Description
ESTALE	1134	046E	STALE	70	0046	Stale NFS file handle

Chapter 21. Reason codes

This chapter lists reason codes that are returned by the z/OS Network File System server and the z/OS Network File System client.

Network File System reason codes are generic eight-digit hexadecimal codes that provide an indication of the problem location. They appear in the format *6Exxyyyy*, where:

6E indicates that this is an NFS reason code.

xx is a two-digit hexidecimal id of the module where the reason code was generated. 0x00 – 0x3F apply to the NFS client and 0x40 – 0xFF apply to the NFS server.

yyyy
is a four-digit hexidecimal line number of the location in the code where the reason code was generated.

You can use Table 67 for initial translation of the reason code *reasoncd* information presented in client messages related to parsing errors.

Table 67. Parsing error (when reason code is 6E01xxxx)

Last 4 hex digits of *reasoncd*	**Reason**
F*xxx*	Unknown keyword
11*yy*	Host name
12*yy*	Path name
13*yy*	Keyword **acdirmax**
14*yy*	Keyword **acdirmin**
15*yy*	Keyword **acregmax**
16*yy*	Keyword **acregmin**
17*yy*	Keyword **cln_ccsid**
18*yy*	Keyword **srv_ccsid**
19*yy*	Keyword **hard**
1A*yy*	Keyword **soft**
1B*yy*	Keyword **retrans**
1C*yy*	Keyword **timeo**
1D*yy*	Keyword **wsize**
1E*yy*	Keyword **rsize**
1F*yy*	Keyword **retry**
2A*yy*	Keyword **vers**
2B*yy*	Keyword **proto**
21*yy*	Keyword **biod**
22*yy*	Keyword **bufhigh**
23*yy*	Keyword **delaywrite**
24*yy*	Keyword **readahead**

© Copyright IBM Corp. 1991, 2006

Table 67. Parsing error (when reason code is 6E01xxxx) (continued)

Last 4 hex digits of *reasoncd*	**Reason**
25*yy*	Keyword **attrcaching**
26*yy*	Keyword **datacaching**
27*yy*	Keyword **dynamicsizeadj**
28*yy*	Keyword **delim**
29*yy*	Keyword **xlat**

Notes:

xxx	Offset to the beginning of the **mount** parameter of the bad keyword.
yy	See Table 68 for more details.

You can use Table 68 for any additional translation of the reason code *reasoncd* information presented in client messages related to parsing errors.

Table 68. Parsing error (when reason code is from 6E0111yy to 6E0129yy)

Last 2 hex digits of *reasoncd*	**Reason**
01	Null host name or null path name
02	Blank detected
03	Incorrect member name in the path name
04	Missing double quotation mark
05	No member name found
06	Missing left parenthesis
07	Incorrect number
08	Number is larger than 2 GB
09	Incorrect multiplier; must be K, M, or G
0A	Missing right parenthesis
0B	The specified number is not within the allowable range
0C	Incorrect keyword parameter value
0D	Mutually exclusive keyword/option
0E	Keyword is not allowed in the mount option
0F	Keyword is not allowed in the installation parameter

Table 69 contains the rest of the reason code *reasoncd* information that is presented in client messages, including the return codes *retc*.

Table 69. Parsing error (when reason code is 6Exxyyyy')

xx **of** *reasoncd*	**Module**	*yyyy* **of** *reasoncd*	**Description**
02	TCP/IP common error	0001	**clntudp_create( )** failed
02	TCP/IP common error	0002	Server NFS port is not 2049
02	TCP/IP common error	0003	**authunix_create( )** failed
02	TCP/IP common error	0004	**clnt_control( )** timeout failed
02	TCP/IP common error	0005	**clnt_control( )** total timeout failed
02	TCP/IP common error	0006	**clnttcp_create( )** failed
02	TCP/IP common error	0007	**clntudp_bufcreate( )** failed
02	TCP/IP common error	*lnum*	**clnt_call( )** timeout (*retc*=0467h, ETIMEDOUT)
02	TCP/IP common error	*lnum*	**clnt_call( )** EINTR (*retc*=0078h, EINTR)
03	TCP/IP error	*yyyy*	TCP error - **clnt_control( )** failed
03	TCP/IP error	*yyyy*	TCP error - **authunix_create( )** failed
03	TCP/IP error	*yyyy*	TCP error - **clnt_call( )** failed
04	TCP/IP error	*yyyy*	Authentication error - **authunix_create( )** failed
05	NFS protocol error	0001	0001, not owner (*retc*=0088h, EPERM)
05	NFS protocol error	0002	00002, no such file or directory (*retc*=0081h, ENOENT)
05	NFS protocol error	0005	00005, I/O error (*retc*=007Ah, EIO)
05	NFS protocol error	0006	00006, no such device or address (*retc*=008Ah, ENXIO)
05	NFS protocol error	000D	00013, permission denied (*retc*=006Fh, EACCES)
05	NFS protocol error	0011	00017, file/dir exists (*retc*=0075h, EEXIST)
05	NFS protocol error	0012	00018, cross-device link (*retc*=0090h, EXDEV)
05	NFS protocol error	0013	00019, no such device (*retc*=0080h, ENODEV)
05	NFS protocol error	0014	00020, not a directory (*retc*=0087h, ENOTDIR)
05	NFS protocol error	0015	00021, is a directory (*retc*=0078h, EISDIR)
05	NFS protocol error	0016	00022, incorrect arguments (*retc*=0079h, EIVAL)
05	NFS protocol error	001B	00027, file too large (*retc*=0077h, EFBIG)
05	NFS protocol error	001C	00028, no space left on device (*retc*=0085h, ENOSPC)
05	NFS protocol error	001E	00030, Read-only file system (*retc*=008Dh, EROFS)
05	NFS protocol error	001F	00031, too many links (*retc*=007Dh, EMLINK)
05	NFS protocol error	003F	00063, file name too long (*retc*=007Eh, ENAMETOOLONG)
05	NFS protocol error	0042	00066, directory not empty (*retc*=0088h, ENOTEMPTY)
05	NFS protocol error	0045	00069, disk quota exceeded (*retc*=046Dh, EDQUOT)
05	NFS protocol error	0046	00070, stale file handle (*retc*=046Eh, ESTALE)
05	NFS protocol error	0047	00071, too many levels of remote (*retc*=046Fh, EREMOTE)
05	NFS protocol error	2711	10001, bad file descriptor (*retc*=0071h, EBADF)
05	NFS protocol error	2712	10002, not sync (*retc*=0071h, EBADF)

Table 69. Parsing error (when reason code is 6Exxyyyy') (continued)

xx **of** *reasoncd*	**Module**	*yyyy* **of** *reasoncd*	**Description**
05	NFS protocol error	2713	10003, bad cookie (*retc*=0076h, EFAULT)
05	NFS protocol error	2714	10004, operation not supported (*retc*=0086h, ENOSYS)
05	NFS protocol error	2715	10005, buffer too small (*retc*=0462h, ENOBUFS)
05	NFS protocol error	2716	10006, server fault (*retc*=007Ah, EIO)
05	NFS protocol error	2717	10007, bad type (*retc*=008Ah, ENXIO)
05	NFS protocol error	2718	10008, jukebox (*retc*=0070h, EAGAIN)
06 - 7F			Contact the IBM Support Center
80 - FF		UNIX System Services (USS) reason code	See *z/OS UNIX System Services Messages and Codes*, SA22-7807

Appendix A. File size value for MVS data sets

Many NFS procedures (such as nfs_lookup and nfs_getattr) in the NFS protocol require the file size to be returned. This appendix explains some performance and accuracy considerations in obtaining the file size value for MVS data sets. For z/OS UNIX files, the file size is directly available from the underlying file system because it is saved as part of the file metadata.

The meaning of the file size value that is returned by the NFS and how fast the file size is returned depends on the following conditions:

- Whether you use text or binary processing mode
- The type of MVS data set being accessed
- If the data set is system-managed
- If **fastfilesize** processing is used

Storage of the file size value

How the file size value is stored affects how quickly files are accessed and depends on the type of MVS data set used.

System-managed PS, VSAM, and PDSE data sets

Text and binary file size are saved on non-volatile storage (DASD) for quick access and maintained by the server for these data set types:

- Physical sequential (including striped)
- VSAM ESDS
- VSAM KSDS
- VSAM RRDS
- PDSE members

These data sets must be SMS managed. When the NFS accesses a data set for the first time, it performs a read-for-size to get the text or binary file size and stores this value on DASD. Subsequent file size requests from clients do not cause the server to read for size, thus improving performance. However, when the data set is modified outside the server by a non-NFS application (for example, by the TSO/E editor), the stored file size could be incorrect. When the data set is accessed again by the server, read-for size must be done to determine the correct file size.

Migrated system-managed data sets

z/OS DFSMS allows data set attribute accessibility for SMS managed data sets, without having to recall the data set if the data set is migrated under DFSMS/MVS V1R3 or later. Supported data set types are SMS managed PS, VSAM ESDS, VSAM KSDS, VSAM RRDS, PDS, and PDSE. Migrated PDS/PDSE members are not supported.

The z/OS NFS server is able to obtain the attributes of a supported SMS managed migrated data set without recalling the data set. Attributes such as the record format and file size are saved to DASD. Subsequent file size requests do not cause a recall of the supported SMS managed migrated data set, thus improving performance. However, when the data set is modified outside the server by a

© Copyright IBM Corp. 1991, 2006

non-Network File System application (for example, by the TSO/E editor) before it was migrated, the stored file size could be incorrect. When the data set is accessed again by the server, a recall must be done to determine the correct file size.

Non-system-managed, PDS, and direct access data sets

The file size value for non-system-managed data sets, PDS members, and direct access (DA) data sets is cached in virtual storage until timeout but not written to DASD. Therefore, for these types of MVS data sets, the file size value is regenerated after the file is closed or after the server is restarted.

How the file size value is generated

When a file is first accessed (for example with **ls -l** or **dir**), usually the entire file is read to determine its size, except for **recfm(f)** or **recfm(fbs)** where the binary size can be computed without reading the file. If the file is a system-managed PS, VSAM, or PDSE member, both binary and text file sizes are stored on DASD, so that subsequent file size requests do not require the file to be read.

Binary file size can be quickly generated by using **recfm(f)** or **recfm(fbs)** to specify a fixed-length record format for the MVS data set. With this format type, the server pads the last logical record with binary zeros in binary mode processing, because MVS always expects complete logical records. If the application tolerates these zeros, using **recfm(f)** or **recfm(fbs)** allows the binary size to be computed quickly because the number of bytes can be computed from the number of blocks, which is stored by MVS.

If you need the exact file size and are using binary mode processing, map it to a variable-format, sequential data set on DASD so that the NFS does not need to pad a partially filled last MVS logical record to a record boundary.

For reading small files or the beginning of files, the read-for-size might not add any processing time. As the file is being read for size, the beginning of the file is stored in the buffers set aside by the **maxrdforszleft** site attribute, until the buffers are full. When the application reads the beginning of the file, this read is fast because it reads directly from the buffer.

MVS stores the number of blocks (rather than the number of bytes) in an MVS file. For most files, therefore, without reading the entire file, the NFS can only give an estimate of the number of bytes in the file, not the exact number of bytes in the file. Even when the server could get the exact byte count without reading the file, the file size could change depending on the file's processing attributes.

For example, selecting **text** mode processing introduces line terminators such as **lf**, **crlf**, or **\n** into the file, thus changing the perceived size of the file. As another example, suppose you select text mode processing with blank stripping enabled on a fixed-length record format file. That causes the server to remove trailing blanks from each record, again changing the perceived size of the file. In these examples, when you first request a file, the server must read the entire file to determine its exact size in bytes.

Using fastfilesize to avoid read-for-size

If you can use an approximate file size for a PDS, PDSE, DA, or non-system-managed data set, you can specify the **fastfilesize** attribute to improve performance. With this attribute, the server estimates the size without opening and reading the entire file.

PDS members
For PDS and PDSE members, the fastfilesize attribute gets the file size from SPF statistics if they exist; otherwise, a zero file size is returned.

DA data sets
For direct access (DA) data sets, an approximate file size is calculated based on the device characteristics, the number of disk tracks in use, and the block size of the data set.

VSAM
For non-system-managed VSAM data sets, the estimated size using fastfilesize is zero. Therefore, **cat** or **vi** will not show any data.

The **fastfilesize** attribute speeds up data set access by calculating approximate file sizes during data set access. Use this only when you are browsing through files (using the **ls** UNIX command for example) because some commands (such as **cp** or **copy**) might not work correctly if fastfilesize is set. When reading or modifying a data set, the **nofastfilesize** attribute should be used to ensure accurate results.

Using nofastfilesize

When you use the default, **nofastfilesize** attribute, the NFS reads the entire file or member to get the file size. It stores the file size value in cache until timeout. If the server's default has been changed to **fastfilesize**, you can still use the nofastfilesize attribute to override it.

```
$ ls -l "filename,nofastfilesize"
```

Using this attribute might cause a delay when first accessing very large data sets.

Note: When directly mounting on a fully qualified data set name and nofastfilesize is specified, the server must return the mount size as part of getting the attributes for the mount. This can slow down the completion of the mount command.

Appendix B. Time stamps for MVS data sets

This appendix explains how to obtain time stamps for MVS data sets. For z/OS UNIX files, the time stamp is directly available from the underlying file system because it is saved as part of the file metadata.

UNIX file attributes define the following time stamps:

atime The last time the file was accessed (read).

mtime The last time the file was modified (write).

ctime The last time the file status was changed (chmod).

The NFS handles time stamps differently for these types of data sets:

- System-managed PS data sets and system-managed VSAM data sets
- Direct Access data sets and non-system-managed PS data sets
- Non-system-managed VSAM data sets
- PDS and PDSE members

The z/OS NFS server treats all MVS data sets as belonging to the same file system (having the same file system id). Since not all attributes are supported for all DFSMSdfp access methods, but the NFS protocols track attribute support at the file system level, the z/OS NFS server reports to the client that **atime** and **mtime** are supported for all MVS data sets. The z/OS NFS server will generate values as described in this appendix.

The **metatime** attribute which is new in NFS version 4 is not supported by any DFSMSdfp access method, nor by the z/OS NFS server. The new NFS version 4 **change** attribute will be based on **mtime** values as described in this appendix.

Time stamps for system-managed VSAM and PS data sets

For system-managed PS data sets and system-managed VSAM data sets, **atime** and **mtime** are fully maintained, and the **ctime** value is set to the **mtime** value.

Time stamps for non-system-managed PS and DA data sets

For non-system-managed physical sequential (PS) and direct access (DA) data sets, consider the following conditions:

- How time stamps are stored.
- The requirements of your workstation programs.
- The type of multiple virtual system (MVS) data set used to store the file.

Storage of time stamps

For non-system-managed physical sequential (PS) and direct access (DA) data sets, the Network File System temporarily stores the time stamps in virtual storage, but not on direct access storage device (DASD). These cached attributes are purged when the file times out and closes or when the server is restarted. When the file is accessed again, the time stamps are regenerated.

© Copyright IBM Corp. 1991, 2006

Client program requirements

Some workstation-based utilities (such as **make**) rely on date and time stamps to determine whether to recompile. For example, **make** checks the update time of the object file with the source file and recompiles if the source has been updated. Before storing these types of files using the MVS server, examine them to ensure that these attributes are unimportant. In an environment which relies on such utilities, use system-managed PS data sets.

Generating time stamps

This is how NFS generates **atime** and **mtime** values for non-system-managed PS and DA data sets from the MVS dates:

```
atime = mtime = reference_date + time_increment
ctime = creation_date + time_increment
```

The *time_increment* value is either the server local time or 23:59 hours. If *reference_date* value or *creation_date* value is equal to the server local date, the server local time is added. Otherwise, a fixed value of 23 hours and 59 minutes is added.

If the *reference_date* value is zero (that is, the file has not yet been referenced), the **atime** value and the **mtime** value are set equal to **ctime**.

Time stamps for non-system-managed VSAM data sets

The time stamps for these types of data sets are set to the current time.

Time stamps for PDSs and PDSEs

An MVS PDS or PDSE data set can act as a UNIX directory, when mounted by an NFS client to the z/OS NFS server. Members of the PDS or PDSE data sets are files within the UNIX directory. When the client accesses the directory, UNIX-format file time stamps are expected for each file on the client side. File time stamps in UNIX format are part of the attributes required by the NFS protocol for NFS client/server communication.

Based on the NFS protocol, the z/OS NFS server generates the following UNIX time stamps to send to the client:

atime
the time when the file data was last accessed

mtime
the time when the file data was last modified

ctime
the time when the attributes of the file were last changed.

The z/OS NFS server converts MVS time stamps to UNIX time stamps (and vice versa) to match NFS protocol requirements. The server uses the following main time stamp sources to generate UNIX time stamps for MVS conventional (legacy) file systems:

- DSCB (data set control block)
- Master Catalog data set attribute extension (AX) cell
- PDSE member attribute extension (AX) cell
- ISPF member statistics
- TOD (current time_of_day on the server side).

Time stamp generation depends on the NFS operation (such as read, write, or setattr) and the type of data set.

TOD is used to set up current times in internal NFS control blocks if needed.

For PDS or PDSE member create/update access with ISPF, some specific additional statistics for the member are maintained by ISPF. They include the creation date and the last modification date and time. The server supports ISPF statistics for compatibility with TSO/ISPF. The server always creates ISPF statistics for new PDS/PDSE members created by NFS clients. For existing PDS/PDSE member updates by the client, the server creates/updates the member ISPF statistics.

The following tables summarize the time stamp sources for NFS operations when obtaining file attributes for PDS (SMS-managed and non SMS-managed) and PDSE data sets and members.

Table 70. Time stamp sources for PDS and PDSE members

		Time stamp sources for members (files)		
Data set type	**ISPF statistics**	**atime**	**mtime**	**ctime**
PDS (note 1)	Available	ISPF_ modification_ date + ISPF_ modification_ time	ISPF_ modification_ date + ISPF_ modification_ time	ISPF_ creation_ date + time_ increment
	Not available	DSCB_ reference_ date + time_ increment	DSCB_ reference_ date + time_ increment	DSCB_ creation_ date + time_ increment
PDSE (note 2)	Available (not used for time generation)	PDSE AX cell	PDSE AX cell	ctime = mtime

Notes:

1. MVS does not maintain time stamps for members of a PDS, only for the PDS data set. MVS creation and reference dates are maintained in the DSCB control block.

   The UNIX time stamps for PDS members are generated from the DSCB creation and reference dates of the PDS data set containing the members, if the time stamps cannot be generated from the member's ISPF statistics.
2. MVS maintains the PDSE member create/change time stamp (mtime) in the PDSE AX cell. The Server uses a FileAccessMethodService (FAMS) call to retrieve/save the member attributes (containing time stamps) from/to the PDSE AX cell.

   For a PDSE data set/member, the server generates the UNIX time stamps obtained from the Catalog AX cell/PDSE AX cell. The server supports PDSE member's ISPF statistics just for compatibility with ISPF but does not return them to the client.

Table 71. Time stamp sources for PDS and PDSE data sets (directories)

	Time stamp sources for data sets (directories)		
Data set type	**atime**	**mtime**	**ctime**
PDS (non-SMS)	DSCB_ reference_ date + time_ increment	DSCB_ reference_ date + time_ increment	DSCB_ creation_ date + time_ increment

Table 71. Time stamp sources for PDS and PDSE data sets (directories) (continued)

PDS (SMS)	catalog AX cell	catalog AX cell	ctime = mtime
PDSE (SMS)	catalog AX cell	catalog AX cell	ctime = mtime

In Table 70 on page 319 and Table 71 on page 319, *time_increment* is either the server local TOD or 23:59 hours. If *reference_date* or *creation_date* is equal to the server local date, the server local TOD is added. Otherwise, a fixed value of 23 hours and 59 minutes is added. If *reference_date* = 0 (that is, the file has not yet been referenced), **atime** and **mtime** are set equal to **ctime**.

MVS maintains PDSE/SMS-managed PDS atime, mtime time stamps in the Catalog AX cell. The Server uses an SVC26 call to retrieve/save the data set attributes (containing time stamps) from/to the Catalog AX cell.

The server keeps all file time stamps current in internal control blocks in readiness to return time stamps when servicing NFS requests issued from NFS clients. The server updates the time stamp sources (in the DSCB, member ISPF statistics, or attribute extension cells) from internal control blocks only during close file operations depending on file timeout expiration values (attrtimeout, readtimeout, writetimeout) as described in section 'Timeout attributes'.

Setting time stamps

NFS clients can issue SETATTR requests to set **atime** and **mtime** for a system-managed PS or VSAM data set. For PDSE members, setting **mtime** is allowed, but setting **atime** is not supported. PDSE member **mtime** is also maintained by PDSE access methods, so it is modified when a TSO/E user modifies the PDSE member.

Appendix C. NFS attributes

This appendix lists NFS attributes and how they are supported on the z/OS NFS server.

NFS file system attributes for MVS data sets

The z/OS NFS server generates MVS-specific values for certain UNIX file system attributes. Table 72, Table 73, and Table 74 illustrate the MVS values that the z/OS NFS server generates.

Table 72. File system values to get dynamic file system information

Value	Description	Conventional MVS Value
tbytes	Total size, in bytes, of the file system	10000000000
fbytes	Amount of free space, in bytes, in the file system	8000000000
abytes	Amount of space, in bytes, available to the user identified by the authentication in the RPC	80000000
tfiles	Total number of file slots in the file system	200000
ffiles	Number of free file slots in the file system	20000
afiles	Number of free file slots that is available to the user corresponding to the authentication information in the RPC	2000
invarsec	Number in seconds for which the file system is not expected to change	0

Table 73. File system values to get static file system information

Value	Description	Conventional MVS Value
rtmax	Maximum number in bytes for the read request supported by the server	65536 (64KB)
rtpref	Preferred size of the read request	32768 (32KB)
rtmult	Suggested multiple for the size read request	4096
wtmax	Maximum size of a write request supported by the server	65536 (64KB)
wtpref	Preferred size of the write request	32768 (32KB)
wtmult	Suggested multiple for the size of a write request	4096
dtpref	Preferred size of the readdir request	8192
maxfilesize	Maximum size of a file on the system	2 ** 63 - 1
time_delta	File time using setattr	(0,1000000)
Properties:	FSF_LINK FSF_SYMLINK FSF_HOMOGENEOUS FSF_CANSETTIME	1 0 1 1

Table 74. File system values to retrieve POSIX information

Value	Description	Conventional MVS Value
linkmax	Maximum number of hard links	1
name_max	Maximum length of a component file name (file name + attributes)	255

© Copyright IBM Corp. 1991, 2006

Table 74. File system values to retrieve POSIX information (continued)

Value	Description	Conventional MVS Value
no_trunc	Server will reject any name that is longer than the name_max	True
chown_restricted	Change either the owner or the group associated with the data set	True
case_insensitive	Server does not distinguish the case when interpreting file names	True
case_preserving	If True, the server file system will preserve the case of a name during a create, mkdir, mknod, symlink, rename, or link	False

NFS file system attributes for z/OS UNIX file systems

For z/OS UNIX files, Table 75 and Table 76 show the file system values that are returned for NFS attributes.

Note: Static UNIX values are for general reference only and may change. Actual values are provided by UNIX System Services or an underlying physical file system such as zFS, TFS, HFS, or NFS client. Refer to these products for current values.

Table 75. File system values to get static file system information

Value	Description	UNIX value
rtmax	Maximum number in bytes for the read request supported by the server	65536 (64KB)
rtpref	Preferred size of the read request	32768 (32KB)
rtmult	Suggested multiple for the size read request	4096
wtmax	Maximum size of a write request supported by the server	65536 (64KB)
wtpref	Preferred size of the write request	32768 (32K)
wtmult	Suggested multiple for the size of a write request	4096
dtpref	Preferred size of the readdir request	8192

Table 76. File system values to retrieve POSIX information

Value	Description	UNIX Value
linkmax	Maximum number of hard links	2 ** 31
name_max	Maximum length of a component file name (file name + attributes)	255
no_trunc	Server will reject any name that is longer than the name_max	True
chown_restricted	Change either the owner or the group associated with the data set	False
case_insensitive	Server does not distinguish the case when interpreting file names	False
case_preserving	If True, the server file system will preserve the case of a name during a create, mkdir, mknod, symlink, rename, or link	True

NFS protocol attributes

Table 77 provides general reference information about the z/OS NFS server's support of NFS protocol attributes. To determine which attributes are communicated between the client and server for the version of the NFS protocol that you are using, see the corresponding RFC at the Internet Engineering Task Force (IETF) home page http://www.ietf.org/ :

- NFS protocol version 2: RFC 1094
- NFS protocol version 3: RFC 1813
- NFS protocol version 4: RFC 3530.

Table 77. NFS Version 4 Attributes

Num	Attribute Name	Supp_attr value (Note 1, 5)	Comments (Note 2)
1	type	1	File object type (File, Directory, Link, etc.) (Note 5) **UNIX value:** at_mode, NFALL_B **MVS value:** NFS generated
2	fh_expire_type	1	Indicates a file handle is persistent or volatile across server restarts. (Note 5) **UNIX value:** NFSV4 fh4_volatile_any **MVS value:** NFSV4 fh4_volatile_any
3	change	1	Server generated value which is updated when an NFS object attribute or content changes. **UNIX value:** at_ctime64, at_ctimemsec **MVS value:** NFS server generated (Notes 4, 5)
4	size	1	Object size in bytes. **UNIX value:** at_size or 512 if an empty directory (Note 5) **MVS value:** (Note 3)
5	link_support	1	Objects file system supports hard links (Note 5) **UNIX value:** fs_nfsprop.fs_fsf_link **MVS value:** 0 (not supported)
6	symlink_support	1	Objects file system supports symbolic links (Note 5) **UNIX value:** fs_nfsprop.fs_fsf_symlink **MVS value:** 0 (not supported)
7	named_attr	1	Object has named attributes (Note 5) **UNIX value:** 0 (not supported) **MVS value:** 0 (not supported)
8	fsid	1	Unique file system id (Note 5) **UNIX value:** at_dev **MVS value:** (Note 8)
9	unique_handles	1	Distinct file handles are guaranteed to represent different objects. **UNIX value:** 0 **MVS value:** 0
10	lease_time	1	Duration of lease in seconds. **UNIX value:** NFS server site attribute. **MVS value:** NFS server site attribute
11	rdattr_error	1	Readattr error during getattr operation.**UNIX value:** error enum **MVS value:** error enum
12	ACL	0	Objects associated access control list **UNIX value:** undefined **MVS value:** undefined
13	aclsupport	0	Types of ACLs supported by the server **UNIX value:** undefined **MVS value:** undefined
14	archive	0	Object has been archived since last modification. **UNIX value:** undefined **MVS value:** undefined
15	cansettime	1	Server is able to change times as specified in a setattr **UNIX value:** fs_nfsprop.fs_fsf_CanSetTime **MVS value:** 1
16	case_insensitive	1	Server file name comparisons are case insensitive. (Note 5) **UNIX value:** pcfgcaseinsensitive **MVS value:** 1
17	case_preserving	1	Filename case is persevered on this file system. **UNIX value:** pcfgcaseonpreserving **MVS value:** 0
18	chown_restricted	1	Changes to file owner or group requires privileged user. (Note 5) **UNIX value:** 0 **MVS value: 1**
19	filehandle	1	Objects associated file handle. (Note 5) **UNIX value:** NFS server generated **MVS value:** NFS server generated
20	fileid	1	A number uniquely identifying the file on this file system. (Notes 5, 7) **UNIX value:** BPXYATTR.AttrFid **MVS value:** NFS server generated

Table 77. NFS Version 4 Attributes (continued)

21	files_avail	1	Available file slots. (Note 5) **UNIX value:** fs_favail **MVS value:** 2,000
22	files_free	1	Free file slots. (Note 5) **UNIX value:** fs_ffree **MVS value:** 20,000
23	files_total	1	Total file slots. (Note 5) **UNIX value:** fs_files **MVS value:** 200,000
24	fs_locations	0	Alternate locations where this file system may be found. **UNIX value:** undefined **MVS value:** undefined
25	hidden	0	Hidden file with respect to the Windows API. **UNIX value:** undefined **MVS value:** undefined
26	homogeneous	1	file system attributes are the same for all objects within the file system. **UNIX value:** fs_nfsprop.fs_fsf_homogeneous **MVS value:**1
27	maxfilesize	1	Maximum file size. (Notes 3, 5) **UNIX value:** fs_maxfilesize **MVS value:** see notes.
28	maxlink	1	Maximum number of links for this object. (Note 5) **UNIX value:** pcfglinkmax **MVS value:** 1
29	maxname	1	Maximum filename size for this object. (Note 5) **UNIX value:** pcfgnamemax **MVS value:** 256
30	maxread	1	Maximum read size for this object. (Note 5) **UNIX value:** 65536**MVS value:** 65536
31	maxwrite	1	Maximum write size for this object. (Note 5) **UNIX value:** 65536 **MVS value:** 65536
32	mimetype	0	MIME body type/subtype for this object. **UNIX value:** undefined **MVS value:** undefined
33	mode	1	Support of UNIX-style mode and permission bits. (Note 5) **UNIX value:** at_mode **MVS value:** 666 or 777
34	no_trunc	1	Return error rather than truncate the filename if name is greater than *maxname*. **UNIX value:** pcfgnotrunc **MVS value:** 1
35	numlinks	1	Number of hard links to this object. (Note 5) **UNIX value:**at_nlink **MVS value:** file = 1, directory = 2
36	owner	1	The owner's string name. (Notes 5, 9) **UNIX value:** see note. **MVS value:** see note.
37	owner_group	1	The owner's group string name. (Note 9) **UNIX value:** see note. **MVS value:** see note.
38	quota_avail_hard	0	Server rejects a write request and provides the additional disk space that can be allocated to this file or directory. **UNIX value:** undefined. **MVS value:** undefined.
39	quota_avail_soft	0	Server reasonably warns on a write request and provides the additional disk space that can be allocated to this file or directory. **UNIX value:** undefined. **MVS value:** undefined.
40	quota_used	0	Amount of disk space used by this file or directory. **UNIX value:** undefined. **MVS value:** undefined.
41	rawdev	UNIX: 1 MVS: 0	Raw device identifier. i.e. UNIX device major/minor node information. **UNIX value:** at_major, at_minor **MVS value:** undefined.
42	space_avail	1	Disk space available to this user. (Notes 3, 5) **UNIX value:** fs_freespace, fs_blocksize **MVS value:** 80,000,000
43	space_free	1	Filesystem free disk space. (Notes 3, 6) **UNIX value:** fs_freespace, fs_blocksize **MVS value:** 8,000,000,000
44	space_total	1	Filesystem total disk space. (Notes 3, 5) **UNIX value:** fs_usedspace, fs_blocksize **MVS value:** 10,000,000,000
45	space_used	1	Allocated bytes for this file object. (Notes 3, 6) **UNIX value:** at_blocksh, at_blocks, at_blksize **MVS value:** see notes.
46	system	0	System file with respect to the Windows API. **UNIX value:** undefined. **MVS value:** undefined.
47	time_access	1	Last read access satisfied by server for this object. (Notes 4, 5, 6) **UNIX value:** at_atime64 **MVS value:** see notes.

Table 77. NFS Version 4 Attributes (continued)

48	time_access_set	1	Setattr operation to set the time of last access for this object. (Notes 4, 5, 6) **UNIX value:** at_atimechg, at_atimeTOD **MVS value:** support is limited, see notes.
49	time_backup	0	Last backup time for this object. **UNIX value:** undefined. **MVS value:** undefined.
50	time_create	UNIX: 1 MVS: 0	Object time of creation. (Notes 4, 6) **UNIX value:** at_createtime64 **MVS value:** undefined.
51	time_delta	1	Smallest useful server time granularity. (Notes 4, 5) **UNIX value:** fs_time_delta_sec, fs_time_delta_ns **MVS value:** 1 second.
52	time_metadata	UNIX: 1 MVS: 0	Time of last meta-data modification to the object. (Note 4) **UNIX value:** at_ctime64,at_ctimemsec **MVS value:** undefined.
53	time_modify	1	Time of last modification to the object. (Notes 4, 6) **UNIX value:** at_mtime64 **MVS value:** see notes.
54	time_modify_set	1	Setattr operation to set the time of last modification for this object. (Notes 4, 5, 6) **UNIX value:** at_mtimechg, at_mtimeTOD **MVS value:** see notes.
55	mounted_on_fileid	UNIX: 1 MVS: 0	Like *fileid*, but if the target filehandle is a file system root the *fileid* of the underlying file system directory is returned. **UNIX value:** undefined. **MVS value:** same as *fileid*.

Notes:

1. Column indicates z/OS NFS server V4 **supp_attr** attribute setting, where 0 = no support, 1 = supported values. The NFS V4 **supp_attr** is a bitmap of attributes requested, returned, or being set in an NFS client request.
2. UNIX value may contain the UNIX system provided macro and field name. For additional information on UNIX values see *z/OS UNIX System Services File System Interface Reference*.
3. See *z/OS DFSMS Using Data Sets* for additional information on MVS file size limits. See Appendix A, "File size value for MVS data sets," on page 313 for additional information on file sizes.
4. For conventional MVS data sets refer to handling of time stamps in Appendix B, "Time stamps for MVS data sets," on page 317.
5. The constant or generated values are being provided by the NFS server for performance or correct operation.
6. For conventional MVS data sets, values represent those in the data set control block DSCB, Catalog locate, PDSE directory, HSM MIC cell, and SMS DATACLASS.
7. The *fileid* of conventional MVS data sets is generated from data set name and member name using **crc32** checksum algorithms.
8. Returned by z/OS UNIX System Services as a result of the NFS server registration request **v_reg()** during server initialization.
9. The z/OS NFS server does not support name@domain owner strings in this release. The string form of numeric UID and GID are returned and accepted for **setattr, verify**, and **nverify**.

Appendix D. NSM (statd) protocol

The NSM (statd) protocol defines network status monitor (NSM) functions for the Network File System. In z/OS V1R7, the network status monitor, along with the network lock manager (NLM) was integrated into the z/OS NFS address space to improve performance and other functions. This integration changed the statd protocol implementations on the z/OS NFS server, as described below.

Using supported NSM (statd) procedures

The NSM protocol defines six RPC procedures which implement the network status manager. With the integration of NSM into the NFS server address space, procedures 0 through 5 act as null procedures and return no results if invoked externally. Procedure 6, however, is fully implemented. The procedures are as follows:

Procedure 0: do nothing
Procedure 0 (NULL) does nothing

Procedure 1: SM_STAT
Dummy call, always return STAT_FAIL

Procedure 2: SM_MON
Monitor a client host

Procedure 3: SM_UNMON
Unmonitor a client host

Procedure 4: SM_UNMON_ALL
Unmonitor all client hosts

Procedure 5: SM_SIMU_CRASH
Simulate a crash

Procedure 6: SM_NOTIFY
NFS server notifies clients that server is restarting, so clients need to reclaim any locks they previously had on server files. This procedure is fully implemented by the z/OS NFS server.

Note: NSM is only active if the NFS server is started with the NLM attribute set, not with NONLM.

© Copyright IBM Corp. 1991, 2006

Appendix E. Sample CHECKLIST data set

The following is a sample CHECKLIST data set.

```
######################################################################
#
#  OS/390 Network File System Server Sample CHECKLIST
#
#     PROPRIETARY V3 STATEMENT=
#      LICENSED MATERIALS - PROPERTY OF IBM
#      THIS MODULE IS "RESTRICTED MATERIALS OF IBM"
#      5647-A01
#      (C) COPYRIGHT IBM CORPORATION 1998
#      SEE IBM COPYRIGHT INSTRUCTIONS
#     END PROPRIETARY V3 STATEMENT
#     (C) Copyright IBM Corp. 1998
#
######################################################################
#
# This data set contains entries for files/directories that are to be
# exempt from SAF checking even though SAF or SAFEXP is specified as
# the security option.  This file is only used when SAF or SAFEXP is
# specified for the particular data type (i.e. MVS data, HFS data,
# data accessed using the public file handle) AND CHECKLIST option is
# specified as the site attribute.  The entries specified here HAVE TO
# MATCH a subsequent mount point OR be the parent of a subsequent mount
# point to allow SAF checking to be bypassed for everything underneath
# that mount point.  If the entry does not match a subsequent mount
# point and it is not the parent of any subsequent mount point then it
# has no effect whatsoever.
#
# For SAFEXP security option, the EXPORTS requirements have to be
# met first before the CHKLIST requirements are examined.  See
# example below.
#
# Proper care must be used to protect this data set.
#
# This data set is read during server startup processing.  Subsequent
# changes to this data set will not take effect until the server is
# restarted.  Errors found in the file are sent to the system log
# and the server terminates.
#
# Statement Syntax
#
# Entries can be up to 1023 characters long. Lines can be continued
# by placing a '+' or '\' at the end of the line. A '#' anywhere in the
# data set indicates a comment that extends to the end of the line.
# No spaces may appear in the name; however, leading blanks
# are ignored in new or continuation lines.
```

© Copyright IBM Corp. 1991, 2006

```
#
#
# Entries are specified by a line of the following
# syntax:
#
# Name
#
# where:
#
# Name = Prefix or file name (which should match a subsequent mount
#        point or be the parent of a subsequent mount point to allow
#        SAF checking to be bypassed for everything underneath that
#        mount point)
#
# Note: Anything after the name on the same line will be ignored.
#
# Example:   Given that security of SAFEXP is specified and the exports
#            list and checklist entries are as follows:
#
#  Exports:                              Checklist:
#  WEBNFS.PDS                            WEBNFS
#  USER1                                 USER1.PUBLIC
#  /hfs/user1/tmp                        /hfs/user1
#  /hfs/u/public                         /hfs/u/public
#  /hfs/bin
##  The following mount requests will fail because it is not in the
#  exports list: (mount points have to be exportable first)
#    WEBNFS
#    /hfs/user1
#
#  The following mount requests will succeed and SAF checking will
#  be bypassed for everything underneath that mount point (i.e. user
#  do not need to do MVSLOGIN for the files/directories underneath
#  the mount point):
#    WEBNFS.PDS
#    USER1.PUBLIC.PDS
#    /hfs/user1/tmp/dir1
#    /hfs/u/public
#
#  The following mount requests will succeed but SAF checking will
#  be enforced:
#    USER1.SOMEPDS
#    /hfs/bin
#
#
# If 'maplower' was specified in the processing attributes, all
# entries are translated to upper case.
#
# If 'nomaplower' was specified in the processing attributes,
# attention must be given to the case of entries.


WEBNFS
/hfs/tmp
```

Appendix F. NFS system server sample attribute table

You can use the contents of the Figure 35 attribute table file as a NFS server sample. The attributes table can be found as GFSAPATT in the NFSSAMP library.

```
######################################################################
#                                                                    #
#  Z/OS Network File System Server Sample Attribute Table            #
#                                                                    #
#     PROPRIETARY STATEMENT=                                         #
#      LICENSED MATERIALS - PROPERTY OF IBM                          #
#      THIS MODULE IS "RESTRICTED MATERIALS OF IBM"                  #
#      5647-A01                                                      #
#      (C) COPYRIGHT IBM CORPORATION 1991, 2004                      #
#      SEE IBM COPYRIGHT INSTRUCTIONS                                #
#     END PROPRIETARY STATEMENT                                      #
#     (C) Copyright IBM Corp. 1991, 1998                             #
#     (C) Copyright SUN Microsystems, Inc &                          #
#     Electronic Data Systems Corp. 1988, 1989                       #
#                                                                    #
######################################################################
#                                                                    #
#
# This is a prototype site defaults attribute file for the
#  Z/OS Network File System Server Sample Attribute Table            #
#
# '#' character starts a comment.  Comments can appear anywhere.
# White space is ignored when parsing the file.

# Default values are illustrated in the examples in this file

#
# Keywords are not case sensitive. 'BLKS' is the same as 'blks' is
# the same as 'BlKs'.

# All time values are in seconds.

######################################################################
# The following are known as data set creation attributes.           #
######################################################################

# SPACE specifies the amount of primary and secondary space allocated
# for a new data set.  The syntax is:
#
#       SPACE(PRIMARY,SECONDARY)
#
# The secondary field is optional (if omitted, the default is taken).
#
# Dimension of allocation is BLKS, TRKS, or CYLS
# RECS is a synonym for BLKS.

space(100,10), blks
```

Figure 35. NFS system server sample attribute table (Part 1 of 17)

© Copyright IBM Corp. 1991, 2006

```
# RLSE specifies that unused space should be released from the data
# set the first time that a new data set is closed.  For slow clients,
# with long pauses between writes, the RLSE attribute will cause space
# to be released from the primary extent prematurely.  Subsequent
# writes will cause secondary space to be allocated.

norlse

# The record format, or RECFM, defines part of the layout of a data
# set: how the records are physically layed out on disk.
#
# Valid RECFM characters are:
#
#       V  - Variable Length Records (LRECL defines maximum size of
#              any record)
#       F  - Fixed Length Records (LRECL defines the actual length of
#              all records)
#       U  - Undefined Length Records
#
# modified by:
#
#       B  - Records are Blocked (BLKSIZE defines the size of the block)
#       S  - Spanned for variable length records
#            Standard format for fixed length records
#       M  - Machine Control Codes
#       A  - ANSI Control Codes
#
# "A" and "S" are mutually exclusive
# "V", "F", and "U" are mutually exclusive
# "S" is not allowed for DSNTYPE(PDS) and DSNTYPE(LIBRARY)
# (refer to DSNTYPE later in this section.)
#
# The BLKSIZE is the size, in bytes, of a physical block on disk.
# BLKSIZE(0) allows the system to choose an optimized block size.
#
# LRECL defines the size, in bytes, of a logical record in the data set.

recfm(vb), blksize(0), lrecl(8196)

# The data set organization can be one of:
#
#       PS              - Physical Sequential
#       DA              - Direct Access
#       INDEXED         - VSAM KSDS data set
#       NONINDEXED      - VSAM ESDS data set
#       NUMBERED        - VSAM RRDS data set
#
# PS is a good organization for NFS usage, and NONINDEXED is the
# corresponding good VSAM data set for NFS (e.g. with AIX client in
# BINARY mode) usage.

dsorg(ps)

# DSNTYPE specifies whether a PDS or a PDSE is to be created when
# the make directory workstation command is issued.
#
# Valid DSNTYPEs are:
#
#       PDS     - Create a Partitioned Data Set.
#       LIBRARY - Create a Partitioned Data Set Extended.

dsntype(pds)
```

Figure 35. NFS system server sample attribute table (Part 2 of 17)

```
# Number of Directory blocks for PDS allocation

dir(27)

# The MGMTCLAS specifies the management class associated with the
# file creation.
#
# The syntax is:
#
#   mgmtclas(mgmt_class_name)

# The VOLUME (or VOL) attribute enables you to specify the volume
# on which to create the specified data set.
#
# The syntax is:
#
#    volume(volser)

# The UNIT attribute enables you to specify the unit on which to
# create the specified data set.
#
# The syntax is:
#
#    unit(unit_name)

# The following attributes are used to control VSAM data set
# creation. They are used only if the DSORG parameter defines
# the data set to be type INDEXED, NONINDEXED or NUMBERED.
#
# Refer to appropriate IBM MVS documentation for a more
# complete description of these and other data set creation
# attributes.

# The KEYS(LENGTH,OFFSET) attribute enables you to define the key
# length and offset for a VSAM INDEXED (KSDS) data set. It is used
# only if DSORG is INDEXED.
#
# Valid range for LENGTH is from 1 to 255.
# Valid range for OFFSET is from 0 to 32760.

keys(64,0)

# The RECORDSIZE(AVERAGE,MAXIMUM) attribute enables a user to define
# the average and maximum record sizes for a VSAM data set.
# These two values must be equal for NUMBERED (RRDS) data sets.
#
# Valid range is from 1 to 32760.

recordsize(512,4K)

# The SPANNED and NONSPANNED attributes define whether VSAM
# records will span control intervals. This option does not affect
# non-VSAM variable length record data sets. Use the 'S' option
# with the RECFM attribute for non-VSAM data sets.

nonspanned

# The SHAREOPTIONS attribute defines the cross region and cross
# system file sharing allowed for a VSAM data set.
#
# Valid range for each argument is from 1 to 4.

shareoptions(1,3)
```

Figure 35. NFS system server sample attribute table (Part 3 of 17)

```
#####################################################################
# The following are known as processing attributes.                 #
#####################################################################

# There are three timeout types: attributes, reads and writes.
# The various timeout values are used by the system to determine
# when to close and deallocate an inactive data set after the last
# "attribute", "read", or "write" operation.
# The WRITETIMEOUT is usually kept short, because WRITE
# operations result in exclusive locking, and you'll want to release
# the data set.  For slow clients, with long pauses between writes,
# you'll want to increase the WRITETIMEOUT value.
#
# Valid range is from "m" to "n"; where "m" is the argument of
# MINTIMEOUT(m) and "n" is the argument of MAXTIMEOUT(n), unless
# NOMAXTIMEOUT is specified.  In that case, "n" is 32767.
#
# xxxxTIMEOUT(n) indicates to deallocate the data set n seconds after
# the "xxxx" operation;
# NOxxxxTIMEOUT indicates not to deallocate the data set after
# the "xxxx" operation;
# Where "xxxx" can be "ATTR", "READ", or "WRITE".
#
# e.g. WRITETIMEOUT(1) indicates to deallocate the data set 1 second
#      after a write.
#      NOWRITETIMEOUT indicates not to deallocate the data set after
#      a write.
#      READTIMEOUT(90) indicates to deallocate the data set 90 seconds
#      after a read if no further activity against it.

attrtimeout(120), readtimeout(90), writetimeout(30)

# Processing may be TEXT or BINARY
#
# BINARY is good for using MVS as a disk farm for PCs and AIX machines
# and offers better performance.
#
# TEXT should be specified if it is necessary to share data sets
# containing textual data with other MVS applications.
#

binary

#
# MAPPED should be specified when a mixed set of data types are
# to be processed on a single mount point.  The determination
# of whether the data is to be processed as text or binary
# depends on the rules established in the specified side file.
# If a file extension cannot be mapped to text or binary, then
# the data will be processed according to what has been specified
# as binary or text at the mount level, and finally, the site
# level.  If binary or text is specified at the file level,
# the specification overrides the MAPPED specification.
#
# The syntax is:
#
# mapped
#
```

Figure 35. NFS system server sample attribute table (Part 4 of 17)

```
The end of line terminators are:
#
#       CR, CRLF, LF, LFCR, or NOEOL.
#
# They define the conversion of records to line terminators in TEXT
# mode.
# LF should be used for AIX clients;
# CRLF should be used for PC clients.
# (set through client mounts appropriately).  When TEXT mode is
# specified, LF is the default.

LF

# BLANKSTRIP or NOBLANKSTRIP affects the processing of trailing
# blanks when reading and writing to Fixed Record data sets with
# text processing enabled.

blankstrip

# The MAPLEADDOT attribute turns on mapping of a file name starting
# with a leading "." from a client to a legal leading "$" for a
# MVS data set name.  NOMAPLEADDOT turns off this mapping.

mapleaddot

# The MAPLOWER attribute tells the server to map file names to
# uppercase when received from the client in an NFS request and
# to translate from uppercase to lowercase when returned to the client.
# Keywords are not case sensitive and are unaffected by this option.
#
# The NOMAPLOWER attribute tells the server NOT to do any translation.
# i.e. the server is neither to map uppercase when received from
# the client nor to translate to lowercase when returned to the client.
# All the entries in the EXPORTS file are case sensitive.
# The client MOUNT request must specify the MVS qualifier with
# the correct case to successfully match the EXPORTS file entry.

maplower

# RETRIEVE tells the server to recall a migrated data set on read/write
# access. NORETRIEVE will force the return of "Device not available"
# error for migrated files.
# RETRIEVE can be coded in the following ways:
#  RETRIEVE         - nowait for recall.                          @09C
#  RETRIEVE(WAIT)   - wait for recall.
#  RETRIEVE(NOWAIT) - nowait for recall.
# For a quiesced HFS file system:                                 @09A
#  RETRIEVE(WAIT) will suspend a request until the file system    @09A
#                 will be unquiesced.                             @09A
#  Others value will force the return of IO error.                @09A
# RETRIEVE is the default

retrieve

# The FASTFILESIZE attribute tells the server to calculate approximate
# file sizes from available catalog information and disk geometries.
# This approximate size may cause problems with client applications
# since the size is probably inaccurate. The NOFASTFILESIZE may
# result in decreased performance because the server may read a data
# set, applying the defined processing attributes, to determine the
# exact size of the data set as viewed by an NFS client.

nofastfilesize
```

Figure 35. NFS system server sample attribute table (Part 5 of 17)

```
# The SETOWNERROOT keyword tells the server to set the user ID in the
# attributes returned to a client for a specified file to 'root' when
# the client is logged on as superuser.  SETOWNERNOBODY tells the
# server to set the user ID in the attributes to 'nobody' (-2).

setownerroot

# You can have the execute bit for plain files on or off by mount
# point. Turn this option on if you plan to store executables
# or shell scripts on the MVS system on a mount by mount basis.
# It should probably always be off in the site file.
#
# EXECUTEBITON will turn on the execute bits (user, group
# and other) for a mount point's files.

executebitoff

# If the installation intends to customize the translation table, @L3XA
# a new DD card, NFSXLAT, is required in the NFSS startup proc.   @L3XA
#                                                                 @L3XA
# //NFSXLAT       DD      DSN=dataset_name,DSP=SHR                @L3XA
#                                                                 @L3XA
# Where dataset_name is the name of PDS or PDSE whose             @L3XA
# members are the translation tables.                             @L3XA
#                                                                 @L3XA
# The XLAT(member) keyword tells the server which member          @L3XA
# the server is to use as the installation default translation    @L3XA
# table. 'member' is the name of a translation table which        @L3XA
# resides in a PDS or PDSE dataset.                               @L3XA
#                                                                 @L3XA
# The syntax is:                                                  @L3XA
#                                                                 @L3XA
# xlat(member)                                                    #@P3C

# FILEEXTMAP or NOFILEEXTMAP affects the file extension mapping   @L59A
# capability.  FILEEXTMAP turns on file extension mapping and     @L59A
# NOFILEEXTMAP turns it off.  This option can be specified at     @L59A
# the file command level.  The default is NOFILEEXTMAP.           @L59A

nofileextmap                                                     #@L59A

# SIDEFILE(dsname) specifies the name of the data set that        @L59A
# contains the rules for file extension mapping purposes.         @L59A
# If a side file name is specified in the attributes data set     @L59A
# then it is the default side file for this NFS server.           @L59A
# A user can also specify another side file name during a MOUNT   @L59A
# operation to be used along with the default.  The mapping rules @L59A
# will first be searched in the side file specified during MOUNT  @L59A
# and then in the default.  To allow file extension mapping a     @L59A
# side file name must be specified either as a default or in the  @L59A
# MOUNT command.  dsname is a fully-qualified MVS data set name   @L59A
# without quotation marks.  SIDEFILE is only specifiable at the   @L59A
# MOUNT level.  See GFSAPMAP for sample mapping side file and     @L59A
# syntax.                                                         @L59A

# TAG specifies the newly created files should receive a file tag.@LCEA
# NOTAG specifies the newly created file should be untagged.      @LCEA
# The default is NOTAG.                                           @LCEA

  notag                                                          #@LCEA
```

Figure 35. NFS system server sample attribute table (Part 6 of 17)

```
# CLN_CCSID(n) specifies the Coded Character Set Identifier(CCSID)@LC1A
# for the remote mounted file system (NFS client) when text is    @LC1A
# being translated.                                               @LC1A
# The default (if specified) is 819 (ISO 8859-1 ASCII).           @LC1A

  cln_ccsid(819)                                                #@LCEC

# SRV_CCSID(n) specifies the Coded Character Set Identifier(CCSID)@LC1A
# for the local file system (Z/OS NFS Server) when                @L6C
# a new file is being created.                                    @LC1A
# The default (if specified) is 1047 (Latin Open System EBCDIC).  @LC1A

  srv_ccsid(1047)                                               #@LCEC

# CONVSERV(technique) - 'technique' specifies technique-search-order
# which Unicode Service will use for specified srv_ccsid(n) and   @LCEA
# cln_ccsid(n) code pages.                                        @LCEA
# 'technique' consists of up 5 technique-characters corresponding @LCEA
# to the 5 available techniques: R, E, C, L and M.                @LCEA
# The default is CONVSERV(LRE).                                   @LCEA

  convserv(LRE)                                                 #@LCEA

# MVSMNT specifies that in NFS version 4, a lookup will be        @P07M
# treated as if a mount procedure were given for the LOOKUP       @P07M
# result object.                                                  @P07M
#                                                                 @P07M
# For any LOOKUPs that do not specify MVSMNT, any Processing      @P07M
# Attributess that may have been provided will be merged          @P07M
# with any that were in effect for the LOOKUP parent directory.   @P07M
#                                                                 @P07M
# For LOOKUPs that *DO* specify MVSMNT, any other Processing      @P07M
# Attributes provided will be merged with the site defaults.      @P07M
# MVSMNT cannot be specified for any LOOKUP where the parent      @P07M
# directory was navigated to by an mount procedure or a result    @P07M
# of an object that was already LOOKUP'ed with MVSMNT.  This      @P07M
# is to ensure that only a client system mount specifies MVSMNT.  @P07M
#                                                                 @P07M
# MVSMNT is not to be specified as a site default attribute.      @P07M

#####################################################################
# The following are known as site attributes.                       #
#####################################################################

# The following are attributes specifiable ONLY in the site
# file (this file).
#
# Some of these values control internal structures and processing
# within the NFS server. Tuning of these values to improve performance
# should be done incrementally and tested.

# SECURITY attribute control the level of security checking.
# The format of the security keyword is security(mvs,hfs,public)  @L5DA
# where:                                                          @L5DA
# mvs - security option for mvs data access                       @L5DA
# hfs - security option for HFS data access                       @L5DA
# public - security option for data access with the public        @L5DA
#          filehandle                                             @L5DA
# The first positional parameter is required and the other two    @L5DA
# are optional.  When the optional parameters are not specified   @L5DA
# they are assigned the same security as the first parameter.     @L5DA
```

Figure 35. NFS system server sample attribute table (Part 7 of 17)

```
# Four options can be chosen from.  They are:
#  NONE    - No security checking is performed.
#  SAF     - SAF checking is performed in line.
#  EXPORTS - EXPORTS file is used to check security.
#  SAFEXP  - Both SAF and EXPORTS file checks are performed.
#
#
# Defaults are SAFEXP for all data accesses.                      @L5DA

security(safexp,safexp,safexp)                                  # @L5DC

# PCNFSD tells the server to start PCNFS server.
# NOPCNFSD tells the server not to start PCNFS server.
# If not specified PCNFSD, default is NOPCNFSD.

nopcnfsd

# LEADSWITCH tells the server to return '/' as the first character
#            in each export entry.
# NOLEADSWITCH tells the server not to return '/' as the first character
#            in each export entry.
# If not specified NOLEADSWITCH, default is LEADSWITCH.

leadswitch

# MINTIMEOUT and MAXTIMEOUT set allowable values for
# ATTRTIMEOUT, READTIMEOUT, and WRITETIMEOUT. Specify NOMAXTIMEOUT
# to allow NOATTRTIMEOUT, NOREADTIMEOUT and NOWRITETIMEOUT specification
# by clients.
#
# Valid range is from 1 to 32767.

mintimeout(1)
nomaxtimeout

# The logout time is the time limit on inactivity for a given user
# on a machine. When the limit is reached, the user is automatically
# logged out. The user must then do another "mvslogin" to restart
# the session. The time value is specified in seconds.
#
# You would probably want to set LOGOUT to the TSO timeout value that
# is defined at your site.
#
# MAXTIMEOUT does not affect the LOGOUT site attribute.
#

logout(1800)                    # 30 minutes (30 * 60)
```

Figure 35. NFS system server sample attribute table (Part 8 of 17)

```
# The readdirtimeout is a new timeout attribute to control the   @03A
# timeout of the readdir cache used by MVS conventional data     @03A
# sets.  The timeout value controls how long before the readdir  @03A
# results saved in cache are discarded.                          @03A
#                                                                @03A
# Valid range is from 1 to 32,767 seconds.                       @04C
#                                                                @03A
# n can go as low as 1 second  but to avoid the possibility of   @04C
# client hanging (because of network delays and staled cache),   @04C
# n is not recommended to be lower than 5 seconds.               @04A
# n may need to be increased if the network is slow and the      @03A
# accessed directory has a lot of entries.                       @03A
# The default readdirtimeout is 30 seconds.                      @03A

readdirtimeout(30)                                          #   @03A

# The RDRVERF attribute tells the server to do cookie verifier   @05A
# checking for NFS version 3 readdir and readdirplus requests.   @05A
#                                                                @05A
# The NORDRVERF attribute tells the server not to do cookie      @05A
# verifier checking for NFS version 3 readdir and readdirplus    @05A
# requests.                                                      @05A
# The default is NORDRVERF.                                      @05A

nordrverf                                                   #   @05A

#
# The NFSTASKS(n,m,o) defines the number of NFS tasks (or         @LA1C
#                     threads) to spawn.                          @LA1C
#                                                                 @LA1A
# If NFSTASKS(n,m) is specified, then the following is true:      @LA1A
#                                                                 @LA1A
# 'n' is the number of subtasks which handle the asynchronous I/O @L3BC
# operations or short blocking operations (the maximum number of  @LA1C
# concurrent NFSS requests).                                      @L3BC
# 'm' is the number of subtasks which handle the long blocking    @L3BC
# operations (the maximum number of concurrent NFSS recall and    @L3BC
# HFS requests.). Increase this value if your server supports     @L3BC
# lots of active recall or HFS clients.                           @L3BC
#
# Valid range for 'n' is from 1 to 24.                             @01C
# Valid range for 'm' is from 1 to 24.                             @01C
# The sum of 'n' plus 'm' must be less than or equal to 25.        @01C
#
# If NFSTASKS(n,m,o) is specified, then the following is true:    @LA1A
#                                                                 @LA1A
# 'n' is the number of subtasks which handle the asynchronous I/O @LA1A
# operations or short blocking operations (the maximum number of  @LA1A
# concurrent NFSS requests).                                      @LA1A
```

Figure 35. NFS system server sample attribute table (Part 9 of 17)

```
# 'm' is the number of subtasks which handle HFS requests.        @LA1A
# Increase this value if your server supports lots of active      @LA1A
# HFS clients.                                                    @LA1A
# 'o' is the number of subtasks which handle the long blocking    @LA1A
# operations (the maximum number of concurrent NFSS recall        @LA1A
# requests.)  Increase this value if your server supports         @LA1A
# lots of active recall.                                          @LA1A
#
# Valid range for 'n' is from 1 to 24.                            @LA1A
# Valid range for 'm' is from 1 to 100.                    @P6C   @LA1A
# Valid range for 'o' is from 1 to 24.                            @LA1A
# The sum of 'n' plus 'o' must be less than or equal to 25.       @LA1A

nfstasks(8,16,8)                                                 # @02C

#                                                                 @L66A
# The MINTASKS(n,m,o) defines the minimum number of NFS tasks     @L66A
# or threads) allowed to run. Tasks may be terminated for reasons @L66A
# such as 878 or 80A abends.                                      @L66A
#                                                                 @L66A
# 'n' is the minimum number of subtasks which handle              @L66A
#     the asynchronous I/O operations or short blocking           @L66A
#     operations. If the number of active 'short' tasks becomes   @L66A
#     less than 'n' the shutdown process of NFS server will start.@L66A
# 'm' is the minimum number of subtasks which handle HFS          @L66A
#     requests. If the number of active HFS tasks becomes less    @L66A
#     than 'm' the shutdown process of NFS server will start.     @L66A
# 'o' is the minimum number of subtasks which handle long         @L66A
#     blocking operations. If the number of active legacy long    @L66A
#     service tasks becomes less than 'o' the shutdown process    @L66A
#     of NFS server will start.                                   @L66A
# If 'n','m' or 'o' are greater than the corresponding values     @L66A
#  in NFSTASKS, or are not specified, they will default to half   @L66A
#  the NFSTASKS values.                                           @L66A
# Valid range for 'n' is from 1 to 24.                            @L66A
# Valid range for 'm' is from 1 to 100.                           @L66A
# Valid range for 'o' is from 1 to 24.                            @L66A
mintasks(4,8,4)                                                 #@L66A
                                                                 @P07C
# The RESTIMEOUT(n,m) defines resource cleanup timer.             @L3LA
# 'n' is the resource retention period for mounts and
# associated resources which will be removed if the mounts
# have been inactive more than 'n' number of hours.               @L3LA
# 'm' is the time of day to do the cleanup work
# for mounts and associated resources which
# have been inactive more than 'n' number of hours.#
```

Figure 35. NFS system server sample attribute table (Part 10 of 17)

```
# The time of day is specified as a 24 hour local time value
# The starting time is at least 24 hours from NFSS start up
#                                                                    @L3LA
# NFSS will appear slow during this cleanup activity.                @L3LA
# Cleanup will be executed under main task, thus preventing          @L3LA
# any additional work from executing.                                @L3LA
# The value of RESTIMEOUT should be set to a value                   @L3LA
# such that this cleanup activity will occur while                   @L3LA
# NFSS is lightly loaded.                                            @L3LA
#                                                                    @L3LA
# Valid range for 'n' is from 1 to 720.                              @L3LA
# If n is set to 0, NFSS will not remove any mount points.           @L3LA
# Note: NFSS keeps all the information for the                       @L3LA
# inactive mount points, thus creating long chains                   @L3LA
# to be searched.                                                    @L3LA
#                                                                    @L3LA
# Valid range for 'm' is from 0 to 23.                               @L3LA

restimeout(48,0)                                                   #@L3LA

# The CACHEWINDOW attribute limits the number of windows to be used
# for each data set in caching client block writes which are received
# out of order.
# The size of each window is determined by the packet size.
# The suggested value is some small multiple of
# the number of BIODs running on an NFS client.
#
# Valid range is from 1 to 256.

cachewindow(112)                                                   # @P5C

# The HFS attribute specifies a new HFS file system prefix to be    @01A
# imbedded in the mount directory path name. The default value of   @01A
# the HFS file system prefix is /hfs. Mount requests received by    @01A
# the Network File System beginning with the HFS file system        @01A
# prefix value are identified as mount requests for z/OS UNIX      @L7EC
# MVS. The HFS file system prefix value is not part of the path     @01A
# name.                                                             @01A
#                                                                   @01A
# Note: The HFS file system must be mounted locally by z/OS UNIX   @L7EC
#       MVS or the client mount fails.                              @01A
# hfs(prefix)                                                      @P07C

hfs(/hfs)                                                          #@P07C

# fileidsize site attribute:                                       @P08C
#                                                                  @P08C
# The fileidsize attribute tells NFS how to handle the size        @P08C
# of fileids.  Fileids can be recognized either as 32-bit          @P08C
# values or 64-bit values.                                         @P08C
#                                                                  @P08C
```

Figure 35. NFS system server sample attribute table (Part 11 of 17)

```
# Client platforms often copy the returned NFS object fileid to   @P08C
# client internal structure fields such as a UNIX inode.   Older  @P08C
# 32-bit applications and operating systems may not support NFS   @P08C
# fileids larger than 32 bits resulting in value too large errors @P08C
# or unexpected client behaviors.                                 @P08C
#                                                                 @P08C
# IBM recommends upgrading older applications and operating       @P08C
# systems to an environment which supports the latest supported   @P08C
# NFS Protocol Version offering the largest supported fileidsize. @P08C
# If upgrading your environment is not possible, the fileidsize   @P08C
# may be set to a lower value to support older applications if    @P08C
# accessed data can be properly recovered or refreshed.           @P08C
#                                                                 @P08C
# Valid values are 32 or 64.                                      @P08C
#                                                                 @P08C
#     NFS Version 2 maximum value in effect is (32) bits          @P08C
#     NFS Version 3 maximum value in effect is (32) bits          @P08C
#     NFS Version 4 maximum value in effect is (64) bits          @P08C
#                                                                 @P08C
# If NFS Version 2 or Version 3 is in effect, the maximum value   @P08C
# of 32 is used.  When the specified fileidsize value             @P08C
# is larger than the maximum value for the NFS Version in effect, @P08C
# the proper maximum value will be used for generating            @P08C
# the NFS fileid attribute.                                       @P08C
#                                                                 @P08C
# Warning:                                                        @P08C
#                                                                 @P08C
# The fileidsize site attribute should be set to the maximum      @P08C
# value to generate the maximum size in unique ids for NFS        @P08C
# objects within a file system.  Setting the value lower than     @P08C
# the allowed maximum (e.g. NFS Version 4) may result in the      @P08C
# same NFS fileid to be returned to the client for                @P08C
# different server objects.  As a result, reducing the            @P08C
# fileidsize site attribute for client                            @P08C
# compatibility may result in data corruption.                    @P08C
#                                                                 @P08C
# The default is fileidsize(64)                                   @P08C
                                                                 #@P08C
fileidsize(64)                                                   #@P08C

# The LOGICALCHE attribute sets the high water mark for all the
# cache windows combined(in bytes).
#
# Valid range is from 1 to 128MB

logicalcache(16M)                                          # @P5C

# The BUFHIGH attribute sets the high water mark, in
# bytes, of data buffers before a buffer reclamation take place.
# A higher number means more caching and probably better read
# performance.
#
# Valid range is from 1 to 128MB.
#
# (If BUFHIGH + LOGICALCACHE is larger than the available storage
#  in the extended private area (implied by the REGION parameter
#  coded in your PROC) at startup, the NFS server will
#  shutdown immediately.)
```

Figure 35. NFS system server sample attribute table (Part 12 of 17)

```
bufhigh(32M)                                                    # @P5C

# On reaching the BUFHIGH high water mark, a percentage of buffers
# is reclaimed for reuse. This is the PERCENTSTEAL. A higher value
# means a reclaim operation is performed less often, but that
# the cached buffers will be significantly trimmed on each reclaim.
# This can result in a poor read performance, because readahead
# buffers may be stolen.
# Lower values result in more frequent reclaim operations, but the
# cached buffers are not significantly trimmed on each reclaim.
## Valid range is from 1 to 99.

percentsteal(20)

# The READAHEADMAX value defines the maximum read ahead (in bytes)
# during read processing.
# This reduces the amount of synchronous physical I/O required for
# NFS read requests in sequential processing.  It also reduces context
# switching overhead on NFS read requests by allowing more read
# requests to be satisfied directly from the main task.
#
# The number is usually set to 2 to 4 times the common block size
# used for file access (which is recommended at 8K for AIX file
# activity).
#
# Valid range is from 1 to 128K

readaheadmax(16K)

# The MAXRDFORSZLEFT attribute defines the number of physical block
# buffers to be remained after a read-for-size operation.
# These buffers are remained to satisfy potential subsequent
# NFS read requests against the same file.
# The buffers remained are subject to trimming during buffer
# steal operations.
#
# Valid range is from 1 to 1024.

maxrdforszleft(32)

# SMF(level,switch) attribute controls the level of smf support  #@08A
# and defines whether to start SMF records collection            #@08A
# at the NFSS startup or not.                                    #@08A
# Level four options can be chosen from. They are:               #@08A
# NONE - No smf records are to be produced.                      #@08A
# USER - User session smf records are to be produced.            #@08A
# FILE - File usage smf records are to be produced.              #@08A
# USERFILE - Both user session and file usage smf records are    #@08A
#            to be produced.                                     #@08A
# Default is NONE.                                               #@08A
# Switch has two options:                                        #@08A
# ON        - activates SMF records collection at the NFSS       #@08A
#             startup                                            #@08A
# OFF       - activation of SMF records collection can be done   #@08A
#             manually by issuing the operator command MODIFY.   #@08A
# Switch is optional parameter. Its default value is OFF.        #@08A
```

Figure 35. NFS system server sample attribute table (Part 13 of 17)

```
#                                                                    #@08A
# The syntax of the SMF attribute is                                 #@08A
# SMF(NONE|USER|FILE|USERFILE,ON|OFF)                                #@08A
# For example,                                                       #@08A
# SMF(USER,ON)                                                       #@08A
# You can use the short form of the SMF attribute:                   #@08A
# SMF(level)                                                         #@08A
# In that case the value of the switch is OFF                        #@08A
smf(none,off)                                                        #@08A

# SFMAX(n) where n specifies the maximum size (in kilobytes) of      @L59A
# allocated storage to hold all of the side files.  n is an          @L59A
# integer from 0 to 2000 (2MB).  The default value is 0 and it       @L59A
# also signifies that no mapping is allowed on this NFS server.      @L59A
# If SFMAX=0, the specification of sidefile in the attributes        @L59A
# data set will cause the server to shut down and the                @L59A
# specification of sidefile in any subsequent MOUNT commands will    @L59A
# cause the mount to fail as mapping is not allowed on this NFS      @L59A
# server.  If the amount of storage specified can not be obtained    @L59A
# during server initialization then the server will shut down        @L59A
# immediately.                                                       @L59A
#                                                                    @L59A
# The default is SFMAX(0).                                           @L59A

sfmax(0)                                                            #@L59A

# PUBLIC(legacy_path,hfs_path) specifies the legacy path             @L53A
# (MVS conventional data) and/or HFS path that is associated with    @L53A
# the public file handle for WebNFS access.  The first path, if      @L53A
# specified, is the legacy path.  The second path is the HFS path    @L53A
# and must start off with the HFS prefix specified in the HFS()      @L53A
# keyword.  If the first path is not there, a comma must precede     @L53A
# the second path.  If the PUBLIC keyword is specified then one      @L53A
# of the paths must be specified.  The PUBLIC keyword must be        @L53A
# specified after the HFS() keyword in this site attribute table.    @L53A
# A LOOKUP request with the public file handle will determine        @L53A
# which of the two paths it is referring to by the pathname that     @L53A
# it comes in with.  An absolute pathname will tell the server       @L53A
# which of the path it is referring to by a match on one of the      @L53A
# paths specified.  A LOOKUP request with a relative pathname        @L53A
# will be taken to be an HFS request if HFS is active (i.e. an       @L53A
# hfs_path has been provided); otherwise, it is treated as a         @L53A
# legacy request.                                                    @L53A
#                                                                    @L53A
# The default is no PUBLIC paths.                                    @L53A
```

Figure 35. NFS system server sample attribute table (Part 14 of 17)

```
# CHECKLIST specifies that the data set pointed to by the CHKLIST @L5DA
# DD in the startup procedure will be scanned.  The data set      @L5DA
# contains a list of entries which should match a subsequent      @L5DA
# mount point or be the parent of a subsequent mount point to     @L5DA
# allow SAF checking to be bypassed for everything underneath     @L5DA
# that mount point.  CHECKLIST is only valid if SAF checking is   @L5DA
# the security option for the particular data access; otherwise,  @L5DA
# it is ignored even if specified.  See GFSAPCHK in NFSSAMP       @L5DA
# library for sample CHKLIST data set.                            @L5DA
# NOCHECKLIST will cause the server not to look at the information@L5DA
# in the CHKLIST data set even if it is there.  This is the       @L5DA
# default.                                                        @L5DA

nochecklist                                                    # @L5DA

# FN_DELIMITER specifies a character to be used instead of a      @L5XA
# comma to delimit the file name from the accompanying attributes.@L5XA
# This will allow those sites which have UNIX data sets containing@L5XA
# commas to copy and store their data on the Z/OS NFS Server.     @L6C
#                                                                 @L5XA
# An example of this would be:                                    @L5XA
#   The attribute data set contains fn_delimiter(;)               @L5XA
#    this changes the default delimiter character from a comma    @L5XA
#    to a semicolon.                                              @L5XA
#                                                                 @L5XA
#   so instead entering ... vi "data_set_name,text,lf"            @L5XA
#                                                                 @L5XA
#   you would enter.........vi "comma,in-name;text,lf"            @L5XA
#                                                                 @L5XA
# This is a new SITE attribute and it's default is a comma        @L5XA
#                                                                 @L5XA
fn_delimiter(,)                                                #@L5XA

# The hfsfbtimeout is a new timeout attribute to control the      @06A
# timeout of the HFS vnode token used by NFS server.              @06A
# The timeout value controls how long before vnode tokens         @06A
# saved in file blocks are released.                              @06A
#                                                                 @06A
# Valid range is from 1 to 32,767 seconds.                        @06A
#                                                                 @06A
# n can go as low as 1 second  but to avoid the possibility of    @06A
# client hanging (because of network delays),                     @06A
# n is not recommended to be lower than 5 seconds.                @06A
# n may need to be increased if the network is slow and the       @06A
# accessed directory has a lot of entries.                        @06A
# The default hfsfbtimeout is 60 seconds.                         @06A
#                                                                 @06A
hfsfbtimeout(60)                                               #@06A
```

Figure 35. NFS system server sample attribute table (Part 15 of 17)

```
# The UPCASE attribute will convert the messages to uppercase.    @07A
# The MIXCASE attribute will display the messages without         @07A
# any modification (as they are).                                 @07A
# The default is upcase.                                          @07A
#                                                                 @07A
upcase                                                           #@07A

# NOREC878 specifies that recovery processing of 878/80A abends  @L66A
# will be turned OFF. That is, if this type of abend occurs,     @L66A
# server will shutdown without recovery.                         @L66A
# It can be used for Debug only.                                 @L66A
# REC878 specifies that recovery processing of 878/80A abends    @L66A
# will be turned ON, and affected tasks will attempt to recover. @L66A
# This is default value.                                         @L66A
rec878                                                           #@L66A

# NLM tells the server to start Lockd/Statd.                     @L74A
# NONLM tells the server not to start Lockd/Statd.               @L74A
# If not specified NLM, default is NONLM.                        @L74A

nonlm                                                            #@L74A

# Leasetime(n)  specifies the lease time for NFSv4 access.  The  @L74A
# lease time is the length of time for which a client is         @L74A
# irrevocably granted a lock.  At the end of a lease period the  @L74A
# lock may be revoked if the lease has not been extended.        @L74A
# The lock must be revoked if a conflicting lock has been        @L74A
# granted after the lease interval.  The lease time is also      @L74A
# used to specify the grace period duration.  The grace period   @L74A
# is the length of time provided for NFSv4 and NLM clients to    @L74A
# reclaim locks following a server crash and restart.            @L74A
#                                                                @L74A
# n is from 5 seconds to 300 seconds.                            @L74A
# The default value is 120 seconds.                              @L74A
# n must be smaller than logout value if logout is not 0.        @L74A

leasetime(120)                                                   #@L74A

# The REMOUNT attribute enables the NFS Server to process NFS     @0AA
# requests after the NFS Server is restarted even though the HFS  @0AA
# file system was remounted with a new HFS file system number     @0AA
# (device number) after its last usage.                          @0AA
# Use of the REMOUNT attribute causes the NFS Server to          @0AA
# automatically access a remounted HFS file system even though it @0AA
# may have been changed prior to remounting.                     @0AA
# Any active client mounts are reestablished.                    @0AA
# The NOREMOUNT attribute enables the NFS Server to fail NFS     @0AA
# requests (with return value NFSERR_STALE) if the HFS file system@0AA
# was remounted with a new HFS file system number (device number) @0AA
# after its last usage.                                          @0AA
# The default is NOREMOUNT.                                      @0AA
# This site attribute applies only to HFS file systems.          @0AA
#                                                                @0AA
```

Figure 35. NFS system server sample attribute table (Part 16 of 17)

```
noremount                                                          #@0AA

# DHCP specifies that NFS Clients using Dynamic IP addresses      @L76A
# are supported.  To function properly, this support requies      @L76A
# that the Clients have static Host Names and that they           @L76A
# dynamically update the DNS server with their IP address         @L76A
# changes.                                                        @L76A
# NODHCP specifies that only NFS Clients using Static IP          @L76A
# addresses are supported.                                        @L76A
# The default is NODHCP.                                          @L76A

nodhcp                                                           #@L76A

# DENYRW specifies that the NFS Version 4 OPEN operation,       @L74M5A
# and the NLM Protocol NLM_Share procedure, should support      @L74M5A
# the Windows Share_Deny value.                                 @L74M5A
# NODENYRW specifies that the NFS Version 4 OPEN operation,     @L74M5A
# and the NLM Protocol NLM_Share procedure, should ignore       @L74M5A
# the Windows Share_Deny value for z/OS UNIX file systems (the    @L7EC
# value will be treated as if DENY_NONE was specified).         @L74M5A
# The default is DENYRW.                                        @L74M5A
                                                                 #@P07C
DENYRW                                                           #@P07C
```

Figure 35. NFS system server sample attribute table (Part 17 of 17)

Appendix G. Sample exports data set

Figure 36 contains an example exports data set. The exports data set can be found as GFSAPEXP in the NFSSAMP library.

```
######################################################################
#                                                                    #
#  z/OS   Network File System Server Sample EXPORTS                  #
#                                                                    #
#     PROPRIETARY STATEMENT=                                         #
#      LICENSED MATERIALS - PROPERTY OF IBM                          #
#      THIS MODULE IS "RESTRICTED MATERIALS OF IBM"                  #
#      5647-A01                                                      #
#      (C) COPYRIGHT IBM CORPORATION 1991, 2003                      #
#      SEE IBM COPYRIGHT INSTRUCTIONS                                #
#     END PROPRIETARY STATEMENT                                      #
#     (C) Copyright IBM Corp. 1991, 1998                             #
#     (C) Copyright SUN Microsystems, Inc &                          #
#     Electronic Data Systems Corp. 1988, 1989                       #
#                                                                    #
######################################################################
#
# This data set contains entries for directories that may be
# exported to Network File System Clients. It is used by the
# server to determine which data sets and prefixes may be
# accessed by a client, and to write protect data sets on the
# server provided that the SECURITY site attribute is set
# to either SECURITY(EXPORTS) or SECURITY(SAFEXP).  This file
# is not used for SECURITY(SAF) or SECURITY(NONE).
#
# The mvsnfs.cntl(exports) data set is read during server startup
# processing. Subsequent changes to this data set will not take
# effect until the server is restarted or the EXPORTFS command
# is issued.  However, changes to the data set do NOT affect the
# mounts processed before the server was restarted or before
# EXPORTFS was issued.
#
# Errors found in the file are sent to the system log, and the server
# continues processing for minor errors such as undefined host
# names.  Termination is brought about if the EXPORTS data set
# cannot be read or if a syntax error exists.  (In the case of the
# EXPORTFS command, these errors will cause the command to be
# ignored and processing will continue with the original exports
# file in effect.)
#
# Statement Syntax
#
# Entries can be up to 4096 characters long. Lines can be continued
# by placing a '+' or '\' at the end of the line. A '#' anywhere in the
# data set indicates a comment that extends to the end of the line.
# No spaces may appear in the keywords; however, leading blanks
# are ignored in new or continuation lines.
#
```

Figure 36. Sample exports data set

© Copyright IBM Corp. 1991, 2006

```
# Wildcard symbols support
#
# The following wildcard symbols can be used to create regular
# expressions in the EXPORTS file.
#
# Symbol          Explanation
#
#  *       Represents any length sequence (can be zero) of any
#          valid characters.
#
# Example  /u/ab*             matches /u/abcde or /u/abfg
#          IBMUSER.HO*.NAME   matches IBMUSER.HOME.NAME or
#                                     IBMUSER.HOLD.NAME
#
#  ?       Represents any (but only one) valid character.
#
# Example  /u/user?           matches /u/user1 or /u/user2
#                             but not /u/user12
#
# [.-.] or
# [...]    Is used to specify one symbol from the region of
#          characters.
#          Characters are regulated in EBCDIC alphabetical order.
#          Open bracket, [, has to be hex value x'AD'
#          Close bracket, ], has to be hex value x'BD'
#
#          NOTE: Hex values x'AD' and x'BD' may represent
#                different characters depending on the
#                code page defined for your installation
#
# Example 1: /u/[a-c]de       matches /u/ade or /u/bde or /u/cde
#            /u/[abd]fg       matches /u/bfg or /u/dfg
#                             but not /u/cfg
# Example 2: Given the set of similar resource names
#
#    /user/abcdef1  /user/abddef1
#
#    /user/abcdef2  /user/abedef2
#
#    /user/abcdef3  /user/abfdef3
#
#    /user/abcdef4
#
# with wildcard symbols support it is possible to export all
# of these names by only one row in the EXPORTS file:
#
#    /user/ab[c-f]def?
# Be aware that with this line other datasets, for example
# /user/abfdef4, are also exported.
#
# MVS system symbols support
#
# The NFS server supports both types of MVS system symbols:
# static and dynamic.
# System-defined static system symbols already have their names
# defined to the system.
# An installation defines a default list of symbols that can be
# adjusted by the system administrator.
#
# Some examples of static MVS system symbols are:
 # &SYSCLONE, &SYSNAME, &SYSPLEX, &SYSR1
# Actual values for static symbols are constant for a running
# system.
#
```

```
# In contrast, dynamic MVS system symbols can change their
# values at any time.
# The NFS server will substitute the actual values for the
# the dynamic symbols at NFS start up and on the EXPORTFS NFS
# operator command.
# These values will be fixed in the NFS exports list until
# the next EXPORTFS command execution even if the value of any
# dynamic symbol changes.
# Some examples of dynamic MVS system symbols are:
# &YYMMDD, &DATE, &HHMMSS, &YR4
#
# When used in the EXPORTS file, MVS system symbols
# (static or dynamic) should be preceded by an "&" sign and
# closed with a ".".
# Also it should not be empty (i.e. "&." is not allowed).
# The name of the MVS system symbol should not exceed 80
# characters.
#
#   Examples            /u/&&SYSNAME.
#                       IBMUSER.&&DATE..REPORT
#
#
#
# Entries for a directory are specified by a line of the following
# syntax:
#
#    directory   -ro
#    directory   -rw=clients
#    directory   -access=clients
#    directory   -sec=sys:krb5:krb5i:krb5p
#
# where:
#
# directory = Prefix or file name
#
# -ro       = Export the directory as read only. If not specified,
#             the directory is exported read/write to every one.
#
# -rw       = Directory exported read/write to specified clients,
#             and read only to everyone else.  Separate clients
#             names by colons.
#
# -access   = Give access only to clients listed. Separate client
#             names by colons.  This can be further qualified with
#             the "ro" keyword or with an "rw" list.
#
# -sec      = Specifies the Kerberos authentication level that
#       clients must have to access individual files and
#       data sets on the z/OS NFS server.
#
# The keywords "ro" and "rw" are mutually exclusive.
#
# If no options are specified, the default value
# allows any client read/write access to the given directory.
#
#
```

```
# If 'nomaplower' is specified in the default attributes,
# attention must be given to the case of entries. High level
# qualifiers specified in the client mount request are not
# translated to upper case.


mvsnfs                   -ro
userid0.mixds            -rw=fsrs001:fslab004:fssun04
userid1                  -rw=fsrs001
userid2.nfs.pds          -access=fsrs001:fssun03
userid3.nfs.ps           -sec=krb5:krb5i:krb5p:sys
userid4.nfs.another.pds  -access=fslab004
/hfs/u/newproducts       -access=fsrs001,ro
```

Appendix H. Sample startup procedures

The following text contains sample startup procedures that you can use to start the z/OS NFS server and client.

Sample z/OS NFS server startup procedures

Figure 37 contains a sample MVSNFS z/OS NFS server startup procedure. This procedure can be found as GFSAPROC in the NFSSAMP library.

```
//MVSNFS  PROC MODULE=GFSAMAIN,PARMS='INFO',
//           SYSNFS=SYS1,SYSLE=SYS1,NFSPRFX=MVSNFS,TCPIP=TCPIP
//*********************************************************************
//*                                                                   *
//*  Z/OS NETWORK FILE SYSTEM SERVER START UP PROC(HDZ11VS)           *
//*                                                                   *
//*PROPRIETARY V3 STATEMENT                                           *
//*LICENSED MATERIALS - PROPERTY OF IBM                               *
//*"RESTRICTED MATERIALS OF IBM"                                      *
//*5694-A01                                                           *
//*(C) COPYRIGHT 1989,2004  IBM CORP.                                 *
//*END PROPRIETARY V3 STATEMENT                                       *
//*                                                                   *
//*********************************************************************
//GFSAMAIN EXEC PGM=&MODULE,
//             PARM='&PARMS',
//             REGION=0M,
//             TIME=1440
//*
//*
//* STEPLIB CONSISTS OF NFSLIBE AND LANGUAGE ENVIRONMENT LIBRARIES
//*         CHANGE THE NAMES AS APPROPRIATE FOR YOUR INSTALLATION
//*  NOTE: EACH OF THESE LIBRARIES MUST BE AUTHORIZED AND
//*        LANGUAGE ENVIRONMENT V1R5M0 OR HIGHER IS REQUIRED
//*
//* &TCPIP..TCPIP.DATA IS TCP/IP DATA FILE
//* &SYSNFS..NFSLIBE IS Z/OS NETWORK FILE SYSTEM SERVER
//*                     TARGET LIBRARY
//* &SYSLE..SCEERUN  IS LANGUAGE ENVIRONMENT RUNTIME LIBRARY
//*
//SYSTCPD  DD   DISP=SHR,DSN=&TCPIP..TCPIP.DATA(TCPDATA)
//STEPLIB  DD   DISP=SHR,DSN=&SYSNFS..NFSLIBE
//         DD   DISP=SHR,DSN=&SYSLE..SCEERUN
//*
//SYSPRINT DD   SYSOUT=*
//*********************************************************************
//*                                                                   *
//* THE SYSMDUMP DATASET NEEDS TO BE PRE-ALLOCATED TO THE             *
//* DCB SPECIFICATIONS REQUIRED BY SYSMDUMP                           *
//*                                                                   *
//*********************************************************************
//SYSMDUMP DD   DISP=SHR,DSN=&NFSPRFX..SYSMDUMP
//OUTPUT   DD   SYSOUT=*
//SYSERR   DD   SYSOUT=*
//*
```

Figure 37. Sample z/OS NFS server startup procedures

© Copyright IBM Corp. 1991, 2006

```
//*********************************************************************
//*                                                                   *
//* NFSLOG1 AND NFSLOG2 DD'S ARE ADDED.                               *
//* THE TWO LOG DATA SETS NEED TO BE PRE-ALLOCATED                    *
//*  AS SEQUENTIAL FILES.                                             *
//* THE DCB WILL BE SET TO (VB,LRECL=137,BLKSIZE=6144)                *
//* ONCE MVSNFS IS STARTED, REGARDLESS OF THE ALLOCATION DCB.         *
//* INITIALLY, WE RECOMMEND SPACE=(CYL,(2,5),RLSE)                    *
//*                                                                   *
//*********************************************************************
//NFSLOG1  DD   DISP=SHR,DSN=&NFSPRFX..LOG1
//NFSLOG2  DD   DISP=SHR,DSN=&NFSPRFX..LOG2
//*
//*  &NFSPRFX..CNTL IS A FB LRECL 80 DATASET (PDS OR PDSE)
//*  SAMPLE NFSATTR CAN BE FOUND IN GFSAPATT IN &NFSPRFX.NFSSAMP
//*  SAMPLE EXPORTS CAN BE FOUND IN GFSAPEXP IN &NFSPRFX.NFSSAMP
//*
//NFSATTR  DD   DISP=SHR,DSN=&NFSPRFX..CNTL(NFSATTR)
//EXPORTS  DD   DISP=SHR,DSN=&NFSPRFX..CNTL(EXPORTS)
//*********************************************************************
//*                                                                   *
//*       FHDBASE AND FHDBASE2 ARE
//*       THE MOUNT FILE HANDLE DATABASES.
//*       THEY NEED TO BE PREALLOCATED
//*       AS EMPTY VSAM KSDS DATA SETS.
//*       THEY WILL BE USED ALTERNATELY.
//*       SAMPLE JCL CAN BE FOUND IN &NFSPRFX.NFSSAMP(GFSAMHDJ).
//*                                                                   *
//*********************************************************************
//FHDBASE  DD   DISP=SHR,DSN=&NFSPRFX..FHDBASE
//FHDBASE2 DD   DISP=SHR,DSN=&NFSPRFX..FHDBASE2
//*********************************************************************
//*                                                                   *
//*       LDBASE AND LDBASE2 ARE
//*       THE LOCKING DATABASES.
//*       THEY NEED TO BE PREALLOCATED
//*       AS EMPTY VSAM KSDS DATA SETS.
//*       THEY WILL BE USED ALTERNATELY.
//*       SAMPLE JCL CAN BE FOUND IN &NFSPRFX.NFSSAMP(GFSAMHDJ).
//*                                                                   *
//*********************************************************************
//LDBASE  DD   DISP=SHR,DSN=&NFSPRFX..LDBASE
//LDBASE2 DD   DISP=SHR,DSN=&NFSPRFX..LDBASE2
//*********************************************************************
//*                                                                   *
//* IF THE INSTALLATION INTENDS TO CUSTOMIZE THE TRANSLATION
//* TABLE, A NEW DD CARD, NFSXLAT, IS REQUIRED IN THE PROC.
//* THE TRANSLATION TABLE DATASET NEEDS TO BE PRE-ALLOCATED
//*  AS PDS OR PDSE DATASET. THE FORMAT OF THE DATASET HAS TO
//*  CONFORM TO THE TRANSLATION TABLE FORMAT SUPPORTED BY
//*  TCP/IP FOR Z/OS.
//*                                                                   *
//*********************************************************************
//*NFSXLAT  DD   DISP=SHR,DSN=&NFSPRFX..XLAT
//*
//*********************************************************************
//*                                                                   *
//* IF THE INSTALLATION INTENDS TO PROVIDE A CHECKLIST DATASET,
//* A NEW DD CARD, CHKLIST, IS REQUIRED IN THE PROC.
//* THE CHECKLIST DATASET NEEDS TO BE PRE-ALLOCATED.
//* THE FORMAT OF THE DATASET HAS TO CONFORM TO THE RULES
//* OUTLINED IN THE SAMPLE CHECKLIST FILE PROVIDED IN GFSAPCHK
//* IN &NFSPRFX.NFSSAMP
//*                                                                   *
```

```
//*********************************************************************
//*CHKLIST  DD   DISP=SHR,DSN=&NFSPRFX..CHKLIST
//*********************************************************************
//* IF IT IS NECESSARY TO USE MVSSNAP() C-FUNCTION
//* A NFSSNAP DD STATEMENT IS REQUIRED.
//* THE NFSSNAP DATA SET NEEDS TO BE PRE-ALLOCATED
//* THE FORMAT OF THE DATA SET IS:
//* DSORG=PS,
//* RECFM=VBA,
//* BLKSIZE=1632,
//* LRECL=125
//*
//* UNCOMMENT THE NFSSNAP DD STATEMENT IN CASE OF USING MVSSNAP()
//*
//*********************************************************************
//*NFSSNAP  DD  DISP=MOD,DSN=&NFSPRFX.NFSSNAP
//NFSSNAP   DD  DUMMY
```

Sample z/OS NFS client startup procedures

Figure 38 contains a sample MVSNFSC z/OS NFS client startup procedure. This procedure can be found as GFSCPROC in the NFSSAMP library.

```
//MVSNFSC PROC SYSNFS=SYS1,SYSLE=SYS1,NFSPRFX=MVSCLNT,TCPIP=TCPIP
//*********************************************************************************
//*
//*  z/OS0 NETWORK FILE SYSTEM CLIENT START UP PROC(HDZ11VC)
//*
//*********************************************************************************
//MVSCLNT  EXEC PGM=BPXVCLNY,
//             REGION=0M,
//             TIME=1440
//*
//*     STEPLIB CONSISTS OF NFSLIBE AND LANGUAGE ENVIRONMENT LIBRARIES
//*              CHANGE THE NAMES AS APPROPRIATE FOR YOUR INSTALLATION
//*       NOTE: EACH OF THESE LIBRARIES MUST BE AUTHORIZED
//*
//* &TCPIP..TCPIP.DATA IS TCP/IP DATA FILE
//* &SYSNFS..NFSLIBE IS z/OS NETWORK FILE SYSTEM CLIENT
//*                     TARGET LIBRARY
//* &SYSLE..SCEERUN IS LANGUAGE ENVIRONMENT RUNTIME LIBRARY
//*
//SYSTCPD  DD   DISP=SHR,DSN=&TCPIP..TCPIP.DATA(TCPDATA)
//STEPLIB  DD   DISP=SHR,DSN=&SYSNFS..NFSLIBE
//         DD   DISP=SHR,DSN=&SYSLE..SCEERUN
//*
//*********************************************************************
//*
//* THE SYSPRINT DD NEEDS TO BE PRE-ALLOCATED
//*  AS A SEQUENTIAL FILE.
//* THE DCB WILL BE SET TO (VB,LRECL=137,BLKSIZE=6144)
//*
//*********************************************************************
//SYSPRINT DD   DISP=SHR,DSN=&NFSPRFX..SYSP
//*********************************************************************
//*
//* THE SYSMDUMP DATASET NEEDS TO BE PRE-ALLOCATED TO THE
//* DCB SPECIFICATIONS REQUIRED BY SYSMDUMP
//*
//*********************************************************************
//SYSMDUMP DD   DISP=SHR,DSN=&NFSPRFX..SYSMDUMP
//*********************************************************************
//*
//* THE SYSOUT, SYSERR AND OUTPUT DD'S NEED TO BE PRE-ALLOCATED
//*  AS A SEQUENTIAL FILES.
//* THE DCB WILL BE SET TO (VB,LRECL=137,BLKSIZE=6144)
//*
```

Figure 38. Sample z/OS NFS client startup procedures

```
//*****************************************************************
//SYSOUT   DD   DISP=SHR,DSN=&NFSPRFX..SYSP
//SYSERR   DD   DISP=SHR,DSN=&NFSPRFX..SYSP
//OUTPUT   DD   DISP=SHR,DSN=&NFSPRFX..SYSP
//SYSIN    DD   DUMMY
//CEEDUMP  DD   DUMMY
//*
//*****************************************************************
//*
//* NFSCMSG1 AND NFSCMSG2 DD'S ARE
//* THE TWO CLIENT LOG DATA SETS NEED TO BE PRE-ALLOCATED
//*  AS SEQUENTIAL FILES.
//* THE DCB WILL BE SET TO (VB,LRECL=137,BLKSIZE=6144)
//* ONCE MVSCLNT IS STARTED, REGARDLESS OF THE ALLOCATION DCB.
//* INITIALLY, WE RECOMMEND SPACE=(CYL,(2,5),RLSE)
//*
//*****************************************************************
//NFSCMSG1 DD   DISP=SHR,DSN=&NFSPRFX..LOG1
//NFSCMSG2 DD   DISP=SHR,DSN=&NFSPRFX..LOG2
```

Appendix I. Retrieval of source code for client enabling commands

Figure 39 shows how you can retrieve the necessary source code to install the client enabling commands on any platform except an AIX, UNIX, or z/OS client. See “Installing the client enabling commands” on page 147 for information on retrieving the necessary source code to install the client enabling commands on AIX or UNIX workstation. The z/OS NFS client source code is installed when the z/OS NFS client and TCP/IP are installed. In this example, **mvshost1** is the name of the z/OS host and **smith** is a z/OS user ID:

```
C> md client         (to create a directory to contain clientcode)
C> cd client
C> ftp mvshost1
Connected to mvshost1
220-FTPSERVE at MVSHOST1, 01:44:24 on 12/02/91
220 Connection closes if idle for more than 5 minutes.
Name (mvshost1:nobody):  smith
<Press ENTER key>
331 Send password please.
Password: password

230 smith is logged on.
ftp> get 'mvsnfs.nfstarb(gfsawaxd)' gfsawaxd.c
200 Port request OK.
125 Sending data set MVSNFS.NFSTARB(GFSAWAXD)
250 Transfer completed successfully.
local: gfsawaxd.c remote: 'mvsnfs.nfstarb(gfsawaxd)'
19516 bytes received in 0.14 seconds (139.4 Kbytes/s)
ftp> get 'mvsnfs.nfstarb(gfsawclt)' gfsawclt.c
200 Port request OK.
125 Sending data set MVSNFS.NFSTARB(GFSAWCLT)
250 Transfer completed successfully.
local: gfsawclt.c remote: 'mvsnfs.nfstarb(gfsawclt)'
17584 bytes received in 0.19 seconds (92.36 Kbytes/s)
ftp> get 'mvsnfs.nfstarb(gfsawlin)' gfsawlin.c
200 Port request OK.
125 Sending data set MVSNFS.NFSTARB(GFSAWLIN)
250 Transfer completed successfully.
local: gfsawlin.c remote: 'mvsnfs.nfstarb(gfsawlin)'
52808 bytes received in 0.25 seconds (211.23 Kbytes/s)
ftp> get 'mvsnfs.nfstarb(gfsawlou)' gfsawlou.c
200 Port request OK.
125 Sending data set MVSNFS.NFSTARB(GFSAWLOU)
250 Transfer completed successfully.
local: gfsawlou.c remote: 'mvsnfs.nfstarb(gfsawlou)'
22468 bytes received in 0.14 seconds (160.49 Kbytes/s)
ftp> get 'mvsnfs.nfstarb(gfsawmcl)' gfsawmcl.c
```

Figure 39. Retrieving source code for client enabling commands (Part 1 of 2)

© Copyright IBM Corp. 1991, 2006

```
200 Port request OK.
125 Sending data set MVSNFS.NFSTARB(GFSAWMCL)
250 Transfer completed successfully.
local: gfsawmcl.c remote: 'mvsnfs.nfstarb(gfsawmcl)'
17302 bytes received in 0.12 seconds (144.18 Kbytes/s)
ftp> get 'mvsnfs.nfstarb(gfsawmou)' gfsawmou.c
200 Port request OK.
125 Sending data set MVSNFS.NFSTARB(GFSAWMOU)
250 Transfer completed successfully.
local: gfsawmou.c remote: 'mvsnfs.nfstarb(gfsawmou)'
29110 bytes received in 0.17 seconds (171.24 Kbytes/s)
ftp> get 'mvsnfs.nfstrab(gfsawrp6)' gfsawrp6.h
200 Port request OK.
125 Sending data set MVSNFS.NFSTARB(GFSAWRP6)
250 Transfer completed successfully.
local: gfsawrp6.h remote: 'mvsnfs.nfstarb(gfsawrp6)'
27880 bytes received in 0.16 seconds (174.25 Kbytes/s)
ftp> get 'mvsnfs.nfstrab(gfsawrs6)' gfsawrs6.h
200 Port request OK.
125 Sending data set MVSNFS.NFSTARB(GFSAWRS6)
250 Transfer completed successfully.
local: gfsawrs6.h remote: 'mvsnfs.nfstarb(gfsawrs6)'
204262 bytes received in 0.94 seconds (216.84 Kbytes/s)
ftp> get 'mvsnfs.nfstarb(gfsawsha)' gfsawsha.c
200 Port request OK.
125 Sending data set MVSNFS.NFSTARB(GFSAWSHA)
250 Transfer completed successfully.
local: gfsawsha.c remote: 'mvsnfs.nfstarb(gfsawsha)'
141122 bytes received in 0.61 seconds (230.97 Kbytes/s)
ftp> get 'mvsnfs.nfstarb(gfsawmnt)' gfsawmnt.h
200 Port request OK.
125 Sending data set MVSNFS.NFSTARB(GFSAWMNT)
250 Transfer completed successfully.
local: gfsawmnt.h remote: 'mvsnfs.nfstarb(gfsawmnt)'
10414 bytes received in 0.09 seconds (115.71 Kbytes/s)
ftp> get 'mvsnfs.nfstarb(gfsawsho)' gfsawsho.h
200 Port request OK.
125 Sending data set MVSNFS.NFSTARB(GFSAWSHO)
250 Transfer completed successfully.
local: gfsawsho.h remote: 'mvsnfs.nfstarb(gfsawsho)'
7052 bytes received in 0.07 seconds (100.74 Kbytes/s)
(Note: Make sure that the tabs key in the gfsawmak makefile will not
 be translated into a space key for the following step.)
ftp> get 'mvsnfs.nfstarb(gfsawmak)' makefile
200 Port request OK.
125 Sending data set MVSNFS.NFSTARB(GFSAWMAK)
250 Transfer completed successfully.
local: makefile remote: 'mvsnfs.nfstarb(gfsawmak)'
8364 bytes received in 0.07 seconds (119.49 Kbytes/s)
ftp> quit
221 Quit command received. Goodbye.
C>
C> (Port commands source code to be compiled in your client
    workstation.)
C> (Set up the path to the "client" directory for the users to use
    the executable code of mvslogin, mvslogout,
    and showattr commands.)
C>  mvslogin mvshost1 smith         (MVSNFS must be operational
Password required                          on host side.)
GFSA973A  Enter MVS password:
GFSA955I  smith logged in ok.
C>
```

Figure 39. Retrieving source code for client enabling commands (Part 2 of 2)

Appendix J. PCNFSD protocol

The PCNFSD protocol enables clients to access z/OS files without issuing the **mvslogin** command. Check with your workstation administrator to see if your clients have PCNFSD support.

Accessing data with PCNFSD

The Network File System creates a unique UID and GID for each client, and the UIDs and GIDs are consistent within a session. For example, the server might return the UID and GID set (100,100) for one user, and (101,100) for the next user.

If the client maintains the UID and GID for the mount point base, you should use the PCNFSD requests rather than the **mvslogin** client enabling command. However, if the client does not maintain the UID and GID for the mount point base, you should use the mvslogin client enabling commands rather than PCNFSD requests.

If you issue the **mount** command, which issues a PCNFSD request (as it is implemented on some platforms), to do the authentication, you should not use the mvslogin client enabling command.

If the client's z/OS user ID, which was authenticated by a PCNFSD request via a mount command, is logged off by the server due to timeout or server restart, then the client needs to issue the **umount** and mount commands to log on again.

See Part 2, "NFS User's Guide," on page 13 for information about how clients use PCNFSD.

Accessing z/OS UNIX files

For z/OS UNIX file access, the user must first have a TSO/E user ID defined to RACF with a z/OS UNIX segment defining the UID and GID. When the PCNFSD authenticate request is received, the Network File System uses the UID and GID associated with the user's RACF profile z/OS UNIX segment, to return in the client's UID and GID credentials. It is recommended that the RACF installation preserve a unique network UID and GID that each user is associated with throughout the network. If there is no z/OS UNIX segment for a user, the client's UID and GID credentials are assigned as presented in the following paragraphs.

If **security(saf)** or **security(safexp)** is specified in the site attributes and the client user does not have a valid z/OS UNIX segment, a system assigned UID and GID is returned in the client user credentials.

If **security(none)** or **security(exports)** is specified in the site attributes, a system assigned UID and GID is returned in the client user credentials.

Starting the PCNFSD server

Specify **pcnfsd** or **nopcnfsd** in the attributes data set to control whether to start the PCNFSD server.

© Copyright IBM Corp. 1991, 2006

Using supported PCNFSD protocols

You can use either version 1 or 2 of the PCNFSD protocol with the Network File System.

Version 1 of the PCNFSD protocol

The PCNFSD Version 1 protocol is specified in *Developer's Specification Protocols for X/Open PC Interworking (PC) NFS*. The z/OS Network File System supports procedures 0 and 1 of the version 1 PCNFSD protocol.

Procedure 0: do nothing

Procedure 0 (NULL) does nothing.

Procedure 1: perform user authentication

The input parameters for this procedure are the z/OS user ID and password. The output parameters are the return code, UID and GID. If the z/OS user ID and password are verified successfully (return code=0), the corresponding UID and GID values are returned.

The caller receives one of these return codes:	• **Return code = 0** - Authentication successful • **Return code = 2** - Authentication failed

Version 2 of the PCNFSD protocol

z/OS NFS supports procedures 0, 1, 12, and 13 of the PCNFSD Version 2 protocol.

Procedure 0: do nothing

Procedure 0 (NULL) does nothing.

Procedure 1: give descriptions

Procedure 1 describes how the NFS supports each procedure of the PCNFSD version 2 Protocol.

The input parameters are the version information and comments from the client platform. The NFS ignores the information in the input parameters.

The output parameters are the version information, comments, number of procedures in the PCNFSD version 2 protocol, and a list of descriptions about how the NFS supports each procedure in this PCNFSD version 2 protocol.

There are three values to represent how the server supports the procedures:	• **-1** - The procedure is not supported • **100** - The procedure gets quick service from the server • **2000** - The procedure takes more time to complete the service from the server

Procedure 12: perform mapping

Procedure 12 performs mapping the UID to user name, GID to group name, or group name to GID depending on the request identifier.

If the mapping is successful, the procedure returns successfully. If the names (or IDs) do not exist, or the caller does not have permission to do the mapping, the procedure returns an error.

Procedure 13: perform user authentication

The input parameters are local host name, encoded user name, encoded password, and comments from the requestor. NFS ignores the local host name and comments from the requestor, and decodes the user name and password in the input parameters. NFS does authentication based on the user name and password.

The output parameters are:

There are only two possible return codes that can be sent back to the caller:	• **Return code = 0** - Authentication successful • **Return code = 2** - Authentication failed
UID and GID	If user name and password are verified successfully (return code = 0), the corresponding UID and GID values are returned.
Number of alternate GIDs	Because NFS does not support alternate GIDs, this field is set to zero.
Pointer to a list of the alternate GIDs	Because NFS does not support alternate GIDs, this field is set to zero.
Logon Home Directory	Because NFS does not support logon home directory, this field is zero.
Default umask	Because NFS does not support default umask, this field is set to zero.
Comments sent to the requestor	If the return code is 2, this comment field contains the reason why the request failed.

Appendix K. SMF C and assembler header macros

The following text contains copies of the SMF C and Assembler header macros for your reference. The SMF C header macro is GFSASSMF and the Assembler header macro is GFSAUSMF. Both header macros can be found in the SYS1.NFSMAC library.

SMF C header macro GFSASSMF

Figure 40 contains information about the SMF C header macro GFSASSMF. The GFSASSMF macro can be found in the SYS1.NFSMAC macro library.

```
/*********************************************************************/ 00050000
/*                                                                   */ 00100000
/*  $MAC(GFSASSMF) COMP(5695DF121): NFSS XI                          */ 00150000
/*                                                                   */ 00200000
/*  MACRO NAME: GFSASSMF.H (SMF.H)                                   */ 00250000
/*                                                                   */ 00300000
/*  DESCRIPTION:  Contains C Language Mapping of SMF type 42         */ 00350000
/*                and the subtype 7 and 8 records.                   */ 00400000
/*                                                                   */ 00450000
/*  STATUS:  Version 1   Release 1                                   */ 00500000
/*                                                                   */ 00550000
/*  COPYRIGHT:                                                       */ 00600000
/*PROPRIETARY V3 STATEMENT                                           */ 00631800
/*LICENSED MATERIALS - PROPERTY OF IBM                               */ 00663600
/*"RESTRICTED MATERIALS OF IBM"                                      */ 00695400
/*5694-A01                                                           */ 00736300
/*(C) COPYRIGHT 1996,2004  IBM CORP.                                 */ 00786300
/*END PROPRIETARY V3 STATEMENT                                       */ 00918000
/*                                                                   */ 00950000
/*  FUNCTION: This header file contains the mapping for SMF          */ 01000000
/*            records in the following format:                       */ 01001500
/*                                                                   */ 01003000
/*            ---------------------------------------                */ 01004500
/*            | Type 42 Header                      |                */ 01006000
/*            ---------------------------------------                */ 01007500
/*            | Subtype 7 or 8 Header               |                */ 01009000
/*            ---------------------------------------                */ 01010500
/*            | Product Section                     |                */ 01012000
/*            ---------------------------------------                */ 01013500
/*            | Subtype 7 or 8 Data                 |                */ 01015000
/*            ---------------------------------------                */ 01016500
/*            | Client Section                      |                */ 01018000
/*            ---------------------------------------                */ 01019500
/*                                                                   */ 01021000
/*  Note: 1) To obtain address of the subtype 7 or 8 header          */ 01022500
/*           add SMF42LEN to the address of the type 42 header       */ 01024000
/*        2) To obtain the address of the product section            */ 01025500
/*           add SMF42OPS to the address of the type 42 header       */ 01027000
```

Figure 40. SMF C header macro GFSASSMF

© Copyright IBM Corp. 1991, 2006

```
/*        3) a) To obtain the address of the subtype 7 data portion  */ 01028500
/*              add SMF42NFO to the address of the type 42 header    */ 01030000
/*           b) To obtain the address of the subtype 8 data portion  */ 01031500
/*              add SMF42NUO to the address of the type 42 header    */ 01033000
/*        4) a) To obtain the address of the subtype 7 client        */ 01034500
/*              section add SMF42FLO to the address of the           */ 01036000
/*              current subtype 7 data portion. There can be more    */ 01037500
/*              than one subtype 7 record in a type 42 record.       */ 01039000
/*              There will always be a subtype 7 and client pair.    */ 01040500
/*           b) To obtain the address of the subtype 8 client        */ 01042000
/*              section add SMF42UCO to the address of the           */ 01043500
/*              type 42 header.                                      */ 01045000
/*                                                                   */ 01046500
/*                                                                   */ 01050000
/*  NOTES:                                                           */ 01100000
/*    DEPENDENCIES:  All changes made to this macro must be          */ 01150000
/*                   reflected in GFSAUSMF. If any of the header     */ 01200000
/*                   sections change, they will also have to be      */ 01250000
/*                   reflected in IGWSMF                             */ 01300000
/*                                                                   */ 01350000
/*  MODULE TYPE:   C header/include file                             */ 01400000
/*    PROCESSOR:   IBM ADC/370 compiler                              */ 01450000
/*    LIBRARY:     NFSLB                                             */ 01500000
/*    CC OPTIONS:  None.                                             */ 01550000
/*    ATTRIBUTES:  None.                                             */ 01600000
/*                                                                   */ 01650000
/*    LINKAGE: #include "GFSASsmf.h"                                 */ 01700000
/*                                                                   */ 01750000
/*                                                                   */ 01800000
/*  EXTERNAL REFERENCES:  (external to this header file)             */ 01850000
/*    DATA AREAS:  None                                              */ 01900000
/*    INCLUDE FILES:  None                                           */ 01950000
/*                                                                   */ 02000000
/* TYPEDEFS DEFINED:                                                 */ 02050000
/*        smf42                                                      */ 02100000
/*        smf427h                                                    */ 02150000
/*        smf428h                                                    */ 02200000
```

```
/*        smf42prd                                                      */ 02250000
/*        smf42s7                                                       */ 02300000
/*        smf42cs                                                       */ 02350000
/*        smf42s8                                                       */ 02400000
/*                                                                      */ 02450000
/* STRUCTs DEFINED:                                                     */ 02500000
/*    None                                                              */ 02550000
/*                                                                      */ 02600000
/* MACROs DEFINED:                                                      */ 02650000
/*    None                                                              */ 02700000
/*                                                                      */ 02750000
/*  CHANGE ACTIVITY:                                                    */ 02800000
/* $L33=NFS,HDZ11NP,940425,SJPLTEM: New macro for DFSMS 1.2/NFSS@L33A*/ 02850000
/* $P1=KA90085,HDZ11NP,940623,SJPLTEM: Add File sys & type values@P1A*/ 02875000
/* $P2=KA90127,HDZ11NP,940726,SJPLTEM: Remove HFS block counters @P2A*/ 02883300
/*                                     Convert to 2 name lengths @P2A*/ 02891600
/* $P3=KA90166,HDZ11NP,940825,SJPLTEM: Add hfs pathname support  @P3A*/ 02895800
/* $P4=KA90179,HDZ11NP,940909,SJPLTEM: Define HFS SMF42FSN & FDN @P4A*/ 02897900
/* $LA5=NFS,HDZ11TS,000521,SJPLJST: AIO C Compiler V2.6 support @LA5A*/ 02898900
/* $L75=NFS,HDZ11VS,031209,IBSVKR:  IPv6 support               @L75A*/ 02899400
/* $01=OA10174,HDZ11US,050311,IBSGYG: for IGWSMF compatibility as@01A*/ 02899600
/*                                    IGWSMF contains SMF42 header01A*/ 02899800
/*********************************************************************/ 02900000
                                                                          03000000
/*-------------------------------------------*/                          03050000
/* Header for SMF record type 42  as in IGWSMF */             /*@01C*/ 03100000
/*-------------------------------------------*/                          03150000
                                                                          03200000
typedef _Packed struct smf42 {                                            03250000
short int      smf42rcl;    /* Record Length                       */  03300000
short int      smf42sgd;    /* Segment Descriptor (RDW) -- 0 if    */  03350000
                            /* record is not spanned               */  03400000
unsigned char  smf42flg;    /* System indicator flags              */  03450000
char           smf42rty;    /* Record type: 42 (X'2A')             */  03500000
int            smf42tme;    /* Time in hundredths of a second      */  03550000
unsigned char  smf42dte[4]; /* Date record written (by SMF)        */  03600000
unsigned char  smf42sid[4]; /* System identification (by SMF)      */  03650000
unsigned char  smf42ssi[4]; /* Subsystem Id                        */  03700000
short int      smf42sty;    /* Record subtype                      */  03750000
short int      smf42nt;     /* Number of triplets (optional)       */  03800000
short int      _filler1;    /* Reserved (optional)                 */  03850000
                                                                          03900000
/*-------------------------------------------*/                          03950000
/* Product section triplet                   */                          04000000
/*-------------------------------------------*/                          04050000
                                                                          04100000
int            smf42ops;    /* Offset to product section           */  04150000
short int      smf42lps;    /* Length of product section           */  04200000
short int      smf42nps;    /* Number of product sections          */  04250000
```

```
/*-------------------------------------------*/                      04350000
/* Header must end on word boundary           */                      04400000
/*-------------------------------------------*/                      04450000
/* unsigned char  smf42end;     1st data section triplet           */  04500000
} SMF42;                                                             04550000
                                                                     04600000
                                                                     04650000
/* Values for field "smf42flg"                                     */  04700000
#define smf42fsi  0x80 /* When set=subsystem id follows system id   */  04750000
#define smf42fsu  0x40 /* When set = subtypes are used             */  04800000
#define smf42fxa  0x04 /* When set = MVS/XA (SMF enters)            */  04850000
#define smf42fs2  0x02 /* When set = VS2 (SMF enters)               */  04900000
#define smf42fs1  0x01 /* When set = VS1 (SMF enters)               */  04950000
                                                                     05000000
#define smf42_len 0x28                                               05050000
                                                                     05100000
                                                                     05150000
/*-------------------------------------------------------------*/    05200000
/* SMF42 subtype 7 header section                              */    05250000
/*  (file timeout statistics)                                  */    05300000
/*-------------------------------------------------------------*/    05350000
                                                                     05400000
typedef struct smf427h {                                             05450000
 int             smf42nfo; /*Offset NFSS file timeout stats section */  05500000
 short int       smf42nfl; /*Length NFSS file timeout stats section */  05550000
 short int       smf42nfn; /*Number NFSS file timeout stats section */  05600000
} SMF427H;                                                           05650000
                                                                     05700000
                                                                     05750000
 /*-------------------------------------------------------------*/   05800000
 /* SMF42 subtype 8 header section                              */   05850000
 /*  (user logout statistics )                                  */   05900000
 /*-------------------------------------------------------------*/   05950000
                                                                     06000000
typedef struct smf428h {                                             06050000
 int             smf42nuo;/* Offset NFSS user session stats section */  06100000
 short int       smf42nul;/* Length NFSS user session stats section */  06150000
 short int       smf42nun;/* Number NFSS user session stats section */  06200000
} SMF428H;                                                           06250000
                                                                     06300000
                                                                     06350000
/*-------------------------------------------*/                      06400000
/* Product Section                           */                      06450000
/*-------------------------------------------*/                      06500000
                                                                     06550000
typedef struct smf42prd {                                            06600000
 char            smf42pdl[8];  /* Product Level                */  06650000
```

```
char            smf42pdn[10]; /* Product Name                     */  06700000
 char            smf42psv;     /* Subtype Version Number       @L75A*/  06733300
 char            filler[21];   /* Reserved                     @L75C*/  06766600
} SMF42PRD;                                                             06800000
                                                                        06850000
                                                                        06900000
/*-------------------------------------------- */                      06950000
/* SMF type 42 subtype 7 file timeout Statistics  */                    07000000
/*-------------------------------------------- */                      07050000
                                                                        07100000
typedef _Packed struct smf42s7 {                              /*@P2C*/  07150000
 int             smf42flo;/*Offset to client identification section.*/  07200000
 short int       smf42fll;/*Length of client identification section.*/  07250000
 char            smf42ffs; /* File system type indicator.           */  07300000
                           /* 01 = HFS                        @P1A*/  07316600
                           /* 02 = MVS                        @P1A*/  07333200
 char            smf42fty; /* File type as defined in NFS protocol. */  07350000
                           /*  0 = Non-file                   @P1A*/  07357100
                           /*  1 = Regular file               @P1A*/  07364200
                           /*  2 = Directory                  @P1A*/  07371300
                           /*  3 = Block special device       @P1A*/  07378400
                           /*  4 = Character special device   @P1A*/  07385500
                           /*  5 = Symbolic link              @P1A*/  07392600
 char            smf42ftm; /* MVS data set type.                */  07400000
                           /*  0 = unknown MVS file type      @P1A*/  07404100
                           /*  1 = Sequential (BSAM) file     @P1A*/  07408200
                           /*  2 = Partitioned (BPAM)         @P1A*/  07412300
                           /*  3 = Direct Access file         @P1A*/  07416400
                           /*  4 = ISAM is not supported      @P1A*/  07420500
                           /*  5 = Virtual Sequential Access  @P1A*/  07424600
                           /*  6 = VSAM Entry Sequenced       @P1A*/  07428700
                           /*  7 = VSAM Relative Record       @P1A*/  07432800
                           /*  8 = VSAM Keyed access          @P1A*/  07436900
                           /*  9 = dummy index level file block @P1A*/  07441000
                           /* 10 = HFS file type              @P1A*/  07445100
 char            s7fill1[3]; /* Reserved                          */  07450000
 int             smf42fsn; /*  File Serial Number, HFS INODE #  @P4C*/  07487500
 int             smf42fdn; /*  Unique device number, HFS file   @P4C*/  07525000
                           /*  system number                  @P4A*/  07562500
 int             smf42fir; /*  Number of I/O blocks read.        */  07600000
 int             smf42fiw; /*  Number of I/O blocks written.     */  07650000
 int             s7fill2;  /*  Reserved                          */  07700000
 u_int64         smf42fbr; /*  Number of bytes read from file. @LA5C*/  07750000
 u_int64         smf42fbw; /*  Number of bytes written to file.@LA5C*/  07800000
 short int       smf42fnl; /* Length of file name.              */  07850000
 char            s7fill3[6];/* Reserved - Dword boudary        @P2A*/  07900000
```

```
/*-----------------------------------------------*/                    07906200
/* This is followed by the file name. It is      */                    07912400
/* a 256 byte character string.             5L75C*/                    07918600
/*                                          @L75D*/                    07924800
/* See the S7RECV0 and S7RECV2 types below. @L75C*/                    07931000
/*                                          @L75D*/                    07937200
/*-----------------------------------------------*/                    07943400
} SMF42S7;                                                             07950000
                                                                       08000000
                                                                       08050000
/*-----------------------------------------------*/                    08100000
/* Client Identification Section. Version 0 @L75C*/                    08150000
/*-----------------------------------------------*/                    08200000
                                                                       08250000
typedef struct smf42cs0 {                                  /*@L75C*/  08300000
 char            smf42cri[8];  /* RACF user ID.                 */    08350000
 char            smf42crg[8]; /* RACF group name.               */    08400000
 char            smf42can[8]; /* Account Number.                */    08450000
 int             smf42cui;    /* User ID at client host (UNIX style)*/ 08500000
 int             smf42cgi;   /* Group ID at client host (UNIX style)*/ 08550000
 int             smf42cip;    /* IP address of client host.     */    08600000
 short  int      smf42chl;    /* Length of client host name     */    08650000
 char            smf42chn[256]; /* Client host name.          @L75C*/  08687500
} SMF42CS0;                                               /*@L75C*/  08693700
                                                                       08699900
                                                                       08706100
/*-----------------------------------------------*/                    08712300
/* Client Identification Section. Version 2 @L75A*/                    08718500
/*-----------------------------------------------*/                    08724700
                                                                       08730900
typedef struct smf42cs2 {                                 /*@L75A*/  08737100
 char            smf42cri[8];   /* RACF user ID.              @L75A*/  08743300
 char            smf42crg[8];   /* RACF group name.           @L75A*/  08749500
 char            smf42can[8];   /* Account Number.            @L75A*/  08755700
 int             smf42cui;      /* User ID at client host     @L75A*/  08761900
 int             smf42cgi;      /* Group ID at client host    @L75A*/  08768100
 char            smf42cip[16];  /* IPv6 address of client host.@L75A*/ 08774300
 short  int      smf42chl;      /* Length of client host name @L75A*/  08780500
 char            smf42chn[256]; /* Client host name.          @L75A*/  08786700
} SMF42CS2;                                               /*@L75A*/  08792900
                                                                       08800000
                                                                       08850000
/*---------------------------------------------------------- */        08900000
/* SMF type 42 subtype 8 user session completion Statistics */        08950000
/*---------------------------------------------------------- */        09000000
                                                                       09050000
typedef struct smf42s8 {                                               09100000
 int             smf42uco;/* Offset client identification section  */ 09150000
 short int       smf42ucl; /*Length client identification section  */ 09200000
 short int       smf42res2; /* Reserved.                       */    09250000
 u_int64         smf42ust; /*Session start time (in STCK format)@LA5C*/ 09300000
 u_int64         smf42uet; /*Session end time (in STCK format)@LA5C*/  09350000
 int             smf42uel; /* Session elapsed time (milliseconds)  */  09400000
 u_long          smf42unr; /* Number of RPC requests processed in  */  09450000
                           /* this session                    */    09500000
 int             smf42ute; /* Total elapsed time of all RPC   */    09550000
                           /* requests processes in this session.  */  09600000
 int             smf42uat; /* Total active time of all RPC    */    09650000
                           /* requests processes in this session.  */  09700000
```

```
u_int64         smf42urn; /*Number of bytes read in from the)@LA5C*/   09750000
                          /* network in this session             */   09800000
 u_int64         smf42uwn; /*Number of bytes written out to the@LA5C*/  09850000
                          /* network in this session             */   09900000
 u_int64         smf42urf; /*Number of bytes read from files on@LA5C*/  09950000
                          /* this session                        */   10100000
 u_int64         smf42uwf; /*Number of bytes written to files in@LA5C*/ 10150000
                          /* this session                        */   10200000
                                                                       10250000
} SMF42S8;                                                             10300000
/*********************************************************************/ 10350000
/* the following typedefs are used internally to construct and   @P2A*/ 10400000
/* free SMF type 42 subtype 7 data chaining elements.            @P2A*/ 10450000
/* elements.                                                     @P2A*/ 10500000
/*                                                               @P2A*/ 10550000
                                                                       10750000
 #define SMF_SHORT_FNAME 256     /* short file name <= 256           */ 10800000
 #define SMF_LONG_FNAME 1023     /* long file name > 256             */ 10850000
 #define SMF_MAX_RECLEN 32390    /* maximum SMF record length        */ 10887500
/* Length of file usage stat., plus short file name, plus client data*/ 10925000
 #define SMF_SHORT_DATALEN      \                                      10950000
    (sizeof(SMF42S7)+SMF_SHORT_FNAME+(GLOBAL->flags.IPv6 ?    /*@L75C*/\10972200
     sizeof(SMF42CS2) : sizeof(SMF42CS0)))                    /*@L75A*/ 10994400
 #define SMF_LONG_DATALEN      \                                       11016600
    (sizeof(SMF42S7)+SMF_LONG_FNAME+(GLOBAL->flags.IPv6 ?     /*@L75C*/\11027700
     sizeof(SMF42CS2) : sizeof(SMF42CS0)))                    /*@L75A*/ 11038800
                                                                       11050000
/* S7ELEM is gotten whenever a file times out, to collect statistics    11060000
   on the file.  It is then chained to the list from NFSGLOBAL, and     11070000
   the chain is processed later to write SMF dataset usage record       11080000
   subtype 7.  This is done for files with file name <= 256.         */ 11090000
 typedef _Packed struct s7elem {         /* element to be chained@P2A*/ 11100000
 _Packed struct s7elem  *next;           /* next element in list @P2A*/ 11150000
         char   *filler;                 /* account for C forcing@P2A*/ 11200000
                                         /* doubleword alignment @P2A*/ 11250000
         SMF42S7 smf42s7;                /* data portion         @P2A*/ 11300000
         char    smfilen[SMF_SHORT_FNAME]; /* 256 byte file name @P2A*/ 11350000
         SMF42CS2 smf42cs;               /* client section V2   @L75C*/ 11400000
  } S7ELEM;                              /*                      @P2A*/ 11450000
                                                                       11500000
/* SMF file usage record. Version 0                            @L75C*/  11533300
typedef _Packed struct s7recV0    {     /* type 42 subtype 7   @L75C*/  11566600
        SMF42    smf42;                 /* type 42 header       @P2A*/  11600000
        SMF427H  smf427h;               /* subtype 7 header     @P2A*/  11650000
        SMF42PRD smf42prd;              /* product section      @P2A*/  11700000
        SMF42S7  smf42s7;               /* data portion         @P2A*/  11750000
        char     smfilen[SMF_SHORT_FNAME]; /* 256 byte path name @P2A*/ 11800000
        SMF42CS0 smf42cs0;              /* client section      @L75C*/  11850000
 } S7RECV0;                             /*                     @L75C*/  11900000
                                                                       11950000
/* SMF file usage record. Version 2                            @L75A*/  11995000
typedef _Packed struct s7recV2 {        /* type 42 subtype 7   @L75A*/  12040000
        SMF42    smf42;                 /* type 42 header      @L75A*/  12085000
        SMF427H  smf427h;               /* subtype 7 header    @L75A*/  12130000
        SMF42PRD smf42prd;              /* product section     @L75A*/  12175000
        SMF42S7  smf42s7;               /* data portion        @L75A*/  12220000
        char     smfilen[SMF_SHORT_FNAME];/* 256 byte path name@L75A*/  12265000
```

```
SMF42CS2 smf42cs2;              /* client section      @L75A*/  12310000
 } S7RECV2;                              /*                    @L75A*/  12355000
                                                                        12400000
/* SMF user session statistics record - subtype 8. Version 0   @L75C*/  12450000
typedef _Packed struct s8recV0   {      /* type 42 subtype 8   @L75C*/  12500000
   SMF42    smf42;                      /* type 42 header       @P2A*/  12550000
   SMF428H  smf428h;                    /* subtype 8 header     @P2A*/  12600000
   SMF42PRD smf42prd;                   /* product section      @P2A*/  12650000
   SMF42S8  smf42s8;                    /* data portion         @P2A*/  12700000
   SMF42CS0 smf42cs0;                   /* client section      @L75C*/  12750000
} S8RECV0;                                                   /*@L75C*/  12800000
                                                                        12850000
/* SMF user session statistics record - subtype 8. Version 2   @L75A*/  12900000
typedef _Packed struct s8recV2   {      /* type 42 subtype 8   @L75A*/  12950000
   SMF42    smf42;                      /* type 42 header      @L75A*/  13000000
   SMF428H  smf428h;                    /* subtype 8 header    @L75A*/  13050000
   SMF42PRD smf42prd;                   /* product section     @L75A*/  13100000
   SMF42S8  smf42s8;                    /* data portion        @L75A*/  13150000
   SMF42CS2 smf42cs2;                   /* client section      @L75A*/  13200000
} S8RECV2;                                                   /*@L75A*/  13250000
```

SMF assembler header macro GFSAUSMF

Figure 41 contains information about the SMF Assembler header macro GFSAUSMF. The GFSAUSMF macro can be found in the SYS1.NFSMAC library.

```
*/*****************************************************************/ 00050000
*/*                                                               */ 00100000
*/* $MAC(GFSAUSMF) COMP(5695DF121): NFSS XI                       */ 00150000
*/*                                                               */ 00200000
*/* MACRO NAME: GFSAUSMF.ASM                                      */ 00250000
*/*                                                               */ 00300000
*/* DESCRIPTION:  Contains ASM Language Mapping of SMF type 42    */ 00350000
*/*               and the subtype 7 and 8 records.                */ 00400000
*/*                                                               */ 00450000
*/* STATUS:  Version 1   Release 1                                */ 00500000
*/*                                                               */ 00550000
*/* COPYRIGHT:                                                    */ 00600000
*/*PROPRIETARY V3 STATEMENT                                       */ 00631800
```

Figure 41. SMF assembler header macro GFSAUSMF

```
*/*LICENSED MATERIALS - PROPERTY OF IBM                                */ 00663600
*/*"RESTRICTED MATERIALS OF IBM"                                       */ 00695400
*/*5694-A01                                                           */ 00736300
*/*(C) COPYRIGHT 1996,2004  IBM CORP.                                 */ 00786300
*/*END PROPRIETARY V3 STATEMENT                                       */ 00918000
*/*                                                                   */ 00950000
*/* FUNCTION: Contains ASM Language Mapping of SMF type 42            */ 01000000
*/*           and the subtype 7 and 8 records in the following        */ 01003000
*/*           format:                                                 */ 01006000
*/*                                                                   */ 01009000
*/*            ---------------------------------------                */ 01012000
*/*            | Type 42 Header                      |                */ 01015000
*/*            ---------------------------------------                */ 01018000
*/*            | Subtype 7 or 8 Header               |                */ 01021000
*/*            ---------------------------------------                */ 01024000
*/*            | Product Section                     |                */ 01027000
*/*            ---------------------------------------                */ 01030000
*/*            | Subtype 7 or 8 Data                 |                */ 01033000
*/*            ---------------------------------------                */ 01036000
*/*            | Client Section                      |                */ 01039000
*/*            ---------------------------------------                */ 01042000
*/*                                                                   */ 01045000
*/*  Note: 1) To obtain address of the subtype 7 or 8 header          */ 01048000
*/*           add SMF42LEN to the address of the type 42 header       */ 01051000
*/*        2) To obtain the address of the product section            */ 01054000
*/*           add SMF42OPS to the address of the type 42 header       */ 01057000
*/*        3) a) To obtain the address of the subtype 7 data portion  */ 01060000
*/*              add SMF42NFO to the address of the type 42 header    */ 01063000
*/*           b) To obtain the address of the subtype 8 data portion  */ 01066000
*/*              add SMF42NUO to the address of the type 42 header    */ 01069000
*/*        4) a) To obtain the address of the subtype 7 client        */ 01072000
*/*              section add SMF42FLO to the address of the           */ 01075000
*/*              current subtype 7 data portion. There can be more    */ 01078000
*/*              than one subtype 7 record in a type 42 record.       */ 01081000
*/*              There will always be a subtype 7 and client pair.    */ 01084000
*/*           b) To obtain the address of the subtype 8 client        */ 01087000
*/*              section add SMF42UCO to the address of the           */ 01090000
*/*              type 42 header.                                      */ 01093000
*/*                                                                   */ 01096000
*/*                                                                   */ 01100000
*/* NOTES:                                                            */ 01150000
*/*   DEPENDENCIES:  All changes made to this macro must be           */ 01200000
*/*                  reflected in GFSASSMF. If changes are made       */ 01250000
*/*                  to any of the header sections, the change        */ 01300000
*/*                  will also have to be made in IGWSMF.             */ 01350000
*/*                                                                   */ 01400000
*/* MACRO:                                                            */ 01450000
*/*   PROCESSOR:   High Level Assembler                               */ 01500000
*/*   LIBRARY:     NFSLB                                              */ 01550000
*/*   ATTRIBUTES:  Include                                            */ 01600000
*/*                                                                   */ 01650000
*/*                                                                   */ 01700000
*/* EXTERNAL REFERENCES:  (external to this header file)              */ 01750000
```

```
*/*   DATA AREAS:     None                                          */ 01800000
*/*   CONTROL BLOCKS: None                                          */ 01850000
*/*   MACROS:         None                                          */ 01900000
*/*                                                                 */ 01950000
*/* CHANGE ACTIVITY:                                                */ 02000000
*/*                                                                 */ 02050000
*/*$L33=NFS,HDZ11NP,940310,SJPLTEM: New macro for DFSMS 1.2/NFSS@L33A*/ 02100000
*/*$P1=KA90085,HDZ11NP,940623,SJPLTEM: Add File sys & type values@P1A*/ 02125000
*/*$P2=KA90127,HDZ11NP,940726,SJPLTEM: Remove HFS block counters @P2A*/ 02137500
*/*$P3=KA90179,HDZ11NP,940909,SJPLTEM: Define HFS SMF42FSN & FDN @P3A*/ 02143700
*/*$L75=NFS,HDZ11VS,031209,IBSVKR:  IPv6 support               @L75A*/ 02146800
*/*$01=OA08867,HDZ11US,040712,IBSVKR: Add Sybtype Version Number @01A*/ 02148400
*/*$02=OA10174,HDZ11US,050311,IBSGYG: for IGWSMF compatibility as@02A*/ 02148900
*/*                                 IGWSMF contains SMF42 header02A*/ 02149400
*/*****************************************************************/ 02150000
         MACRO                                                      02200000
         GFSAUSMF                                                   02250000
*                                                                   02300000
*                                                                   02350000
*********************************************************************    02400000
*     Header for SMF record type 42                               *    02450000
*********************************************************************    02500000
*                                                                   02550000
*                                                                   02600000
SMF42    DSECT                                                      02650000
SMF42BAS DS    0D        SMF42BAS is the basing expr.               02700000
SMF42RCL DS    H         Record Length                              02750000
SMF42SGD DS    H         Segment Descriptor (RDW) -- 0 if record is 02800000
*                        not spanned                                02850000
SMF42FLG DS    0BL1      System indicator flags                     02900000
SMF42FSI EQU   X'80'     When set=subsystem id follows system id    02950000
SMF42FSU EQU   X'40'     When set = subtypes are used               03000000
SMF42FXA EQU   X'04'     When set = MVS/XA (SMF enters)             03050000
SMF42FS2 EQU   X'02'     When set = VS2 (SMF enters)                03100000
SMF42FS1 EQU   X'01'     When set = VS1 (SMF enters)                03150000
         ORG   SMF42FLG+X'00000001'                                 03200000
SMF42RTY DS    X         Record type: 42 (X'2A')                    03250000
SMF42TME DS    FL4       Time in hundredths of a second when record 03287500
*                        was moved to SMF buffer                    03325000
SMF42DTE DS    CL4       Date record written (by SMF)               03362500
SMF42SID DS    CL4       System identification (by SMF)             03400000
SMF42SSI DS    CL4       Subsystem Id                               03450000
SMF42STY DS    HL2       Record subtype                             03500000
SMF42NT  DS    HL2       Number of triplets (optional)              03550000
         DS    HL2       Reserved (optional)                        03600000
*                                                                   03650000
*********************************************************************    03700000
*    Product section triplet                                      *    03750000
*********************************************************************    03800000
*                                                                   03850000
SMF42OPS DS    FL4       Offset to product section                  03900000
SMF42LPS DS    HL2       Length of product section                  03950000
SMF42NPS DS    HL2       Number of product sections                 04000000
*                                                                   04050000
*********************************************************************    04100000
*    Header must end on word boundary                             *    04150000
```

```
*********************************************************************    04200000
*                                                                          04250000
SMF42END DS    0F        1st data section triplet                        04300000
SMF42LEN EQU  *-SMF42                                                    04350000
*                                                                          04400000
*                                                                          04450000
*********************************************************************    04500000
*     Product Section                                               *    04550000
*********************************************************************    04600000
SMF42PSV DS    CL1       Subtype Version Number                @L75A    04833300
         DS    CL21      Reserved                              @L75C    04866600
**                                                                       04877000
*SMF42PRD DSECT                                                          04887400
*SMF42PDL DS    CL8       Product Level                                 04897800
*SMF42PDN DS    CL10      Product Name                                  04908200
*SMF42PSV DS    CL1       Subtype Version Number                 @01A   04918600
*         DS    CL21      Reserved                               @01C   04929000
**                                                                       04939400
*********************************************************************    04950000
*    Product section must end on word boundary                      *    05000000
*********************************************************************    05050000
**                                                                       05100000
*SMF42PEN DS    0F                                                      05150000
*SMF42PLN EQU   *-SMF42PRD Length of product section                    05200000
**                                                                       05250000
**    end of header deletion                                  56@02D   05300000
*********************************************************************    04950000
*    Product section must end on word boundary                      *    05000000
*********************************************************************    05050000
*                                                                          05100000
SMF42PEN DS    0F                                                        05150000
SMF42PLN EQU   *-SMF42PRD Length of product section                      05200000
*                                                                          05250000
*                                                                          05300000
*********************************************************************    05350000
*  SMF42 subtype 7 header section                                  *     05400000
*   (file timeout statistics)                                      *     05450000
*********************************************************************    05500000
*                                                                          05550000
SMF427H  DSECT                                                           05600000
SMF42NFO DS    F         Offset to NFSS file timeout stats section      05650000
SMF42NFL DS    H         Length of NFSS file timeout stats section      05700000
SMF42NFN DS    H         Number of NFSS file timeout stats section      05750000
SMF427HE EQU   *-SMF427H Length of subtype 7 header                      05800000
*                                                                          05850000
*                                                                          05900000
*********************************************************************    05950000
*   SMF42 subtype 8 header section                                 *     06000000
*    (user logout statistics )                                     *     06050000
*********************************************************************    06100000
*                                                                          06150000
SMF428H  DSECT                                                           06200000
SMF42NUO DS    F         Offset to NFSS user session stats section      06250000
SMF42NUL DS    H         Length of NFSS user session stats section      06300000
SMF42NUN DS    H         Number of NFSS user session stats section      06350000
SMF428HE EQU   *-SMF428H Length of subtype 8 header                      06400000
*                                                                          06450000
*                                                                          06500000
*********************************************************************    06550000
*     SMF type 42 subtype 7 file timeout Statistics                 *    06600000
*       (file timeout statistics)                                   *    06650000
*********************************************************************    06700000
*                                                                          06750000
```

```
SMF42S7  DSECT                                                          06800000
SMF42FLO DS    F         Offset to client identification section        06850000
SMF42FLL DS    H         Length of client identification section        06900000
SMF42FFS DS    X         File system type indicator                     06950000
*                            01 = HFS                          @P1A     06966600
*                            02 = MVS                          @P1A     06983200
SMF42FTY DS    X         File type as defined in NFS protocol           07000000
*                            00 = Non-file                     @P1A     07007100
*                            01 = Regular file                 @P1A     07014200
*                            02 = Directory                    @P1A     07021300
*                            03 = Block special device         @P1A     07028400
*                            04 = Character special device     @P1A     07035500
*                            05 = Symbolic link                @P1A     07042600
SMF42FTM DS    X         MVS data set type                              07050000
*                             0 = unknown MVS file type        @P1A     07054100
*                             1 = Sequential (BSAM) file       @P1A     07058200
*                             2 = Partitioned (BPAM)           @P1A     07062300
*                             3 = Direct Access file           @P1A     07066400
*                             4 = ISAM is not supported        @P1A     07070500
*                             5 = Virtual Sequential Access    @P1A     07074600
*                             6 = VSAM Entry Sequenced         @P1A     07078700
*                             7 = VSAM Relative Record         @P1A     07082800
*                             8 = VSAM Keyed access            @P1A     07086900
*                             9 = dummy index level file block @P1A     07091000
*                            10 = HFS file type                @P1A     07095100
         DS    XL3       Reserved                                       07100000
SMF42FSN DS    F         File Serial Number, HFS INODE #       @P3C     07150000
SMF42FDN DS    F         Unique device number HFS file system #@P3C     07200000
SMF42FIR DS    F         Number of I/O blocks read                      07250000
SMF42FIW DS    F         Number of I/O blocks written                   07300000
         DS    F         Reserved                                       07350000
SMF42FBR DS    D         Number of bytes read from file                 07400000
SMF42FBW DS    D         Number of bytes written to file                07450000
SMF42FNL DS    H         Length of file name                            07500000
SMF42FCL DS    0D        C370 ends structure on doubleword boundry      07600000
SMF42FFN DS    0D        File name                          @P2C        07625000
*                                                                       07630000
* The file name is either a 256 byte character string, or a             07635000
* 1023 byte character string. SMF42FNL will contain the length         07640000
*                                                                       07645000
*                        Start of Client Section                        07650000
SMF42F7E EQU   *-SMF42S7 Length of subtype 7 data section               07700000
*                                                                       07712500
SMFSNAME EQU   256       Short file name                                07725000
SMFLNAME EQU   1023      Long file name                                 07737500
*                                                                       07750000
*                                                                       07800000
*********************************************************************   07850000
*  Client Identification Section. Version 0                   @L75C*    07900000
*********************************************************************   07950000
*                                                                       08000000
SMF42CS0 DSECT @L75C                                                    08050000
SMF42CRI DS    CL8       RACF user ID                                   08100000
SMF42CRG DS    CL8       RACF group name                                08150000
SMF42CAN DS    CL8       Account Number                                 08200000
SMF42CUI DS    F         User ID at client host (UNIX style)            08250000
SMF42CGI DS    F         Group ID at client host (UNIX style)           08300000
SMF42CIP DS    F         IP address of client host                      08350000
SMF42CHL DS    H         Length of client host name                     08400000
SMF42CHN DS    CL256     Client host name                     @L75C     08450000
SMF42CSE EQU   *-SMF42CS0 Length of Client Section Version 2   @L75C     08458300
*                                                                       08466600
*                                                                       08474900
*********************************************************************   08483200
*  Client Identification Section. Version 2                  @L75A*     08491500
*********************************************************************   08499800
*                                                                       08508100
SMF42CS2 DSECT                                              @L75A       08516400
```

```
S2F42CRI DS    CL8       RACF user ID                          @L75A    08524700
S2F42CRG DS    CL8       RACF group name                       @L75A    08533000
S2F42CAN DS    CL8       Account Number                        @L75A    08541300
S2F42CUI DS    F         User ID at client host                @L75A    08549600
S2F42CGI DS    F         Group ID at client host               @L75A    08557900
S2F42CIP DS    CL16      IPv6 address of client host           @L75A    08566200
S2F42CHL DS    H         Length of client host name            @L75A    08574500
S2F42CHN DS    CL256     Client host name                      @L75A    08582800
S2F42CSE EQU   *-SMF42CS2 Length of Client Section Version 2   @L75A    08591100
*                                                                       08600000
*                                                                       08650000
*********************************************************************   08700000
*  SMF type 42 subtype 8 user session completion Statistics        *    08750000
*    (user logout statistics )                                     *    08800000
*********************************************************************   08850000
*                                                                       08900000
SMF42S8  DSECT                                                          08950000
SMF42UCO DS    F         Offset to client identification section       09000000
SMF42UCL DS    H         Length of client identification section       09050000
         DS    H         Reserved                                      09100000
SMF42UST DS    D         Session start time (in STCK format)           09150000
SMF42UET DS    D         Session end time (in STCK format)             09200000
SMF42UEL DS    F         Session elapsed time (in milliseconds)        09250000
SMF42UNR DS    F         Number of RPC requests processed in this      09300000
*                        session                                       09350000
SMF42UTE DS    F         Total elapsed time of all RPC requests        09400000
*                        processes in this session                     09450000
SMF42UAT DS    F         Total active time of all RPC requests         09500000
*                        processes in this session                     09550000
SMF42URN DS    D         Number of bytes read in from the network in   09600000
*                        this session                                  09650000
SMF42UWN DS    D         Number of bytes written out to the network    09700000
*                        in this session. */                           09750000
SMF42URF DS    D         Number of bytes read from files on this       09900000
*                        session */                                    09950000
SMF42UWF DS    D         Number of bytes written to files in this      10000000
*                        session                                       10050000
SMF42UCS DS    0F        Start of Client Section                       10100000
SMF42S8E EQU   *-SMF42S8 Length of subtype 8 data section              10150000
         MEND                                                          10200000
```

Appendix L. Capturing diagnostic information using z/OS NFS log data sets and from other components

This appendix describes how to use the z/OS NFS log data sets and how to capture diagnostic information for other components including z/OS UNIX, TCP/IP, and AIX.

Using log data sets

This section describes how to use log data sets for the z/OS NFS server and client.

Server log data sets

The z/OS Network File System server records diagnostic information (messages and debug trace information) in the z/OS NFS server log data sets specified in the NFSLOG1 and NFSLOG2 DD statements of the server's startup procedure.

The z/OS NFS server also records messages and debug trace information in a z/OS component trace buffer if specified. Using a component trace buffer can provide performance improvements compared to the log data sets. For details, see "Using z/OS component tracing" on page 229.

NFSLOG1 is associated with the primary log, while NFSLOG2 is associated with the secondary log. When the server is started, the primary log is used first. When the primary log is full, the logging is automatically switched or toggled to the secondary log and a console message is displayed. The last log buffer is lost when the switch takes place. When the secondary log is also full, the logging is switched back to the primary log and the original primary log content is overwritten.

The number of log records depends on the log level setting and the number of transactions that the server handles. Adjust the space allocation of your z/OS NFS server log data sets according to your installation experience to avoid frequent log switching.

Setting up the z/OS NFS server log data set

Here is an example of setting up the z/OS NFS server log data sets:

```
//NFSLOG1   DD   DISP=SHR,DSN=MVSNFS.PROCESS.LOG1
//NFSLOG2   DD   DISP=SHR,DSN=MVSNFS.PROCESS.LOG2
```

File attributes of the z/OS NFS server log data sets can be set up like this:

```
Organization       PS
Record Format      VB
Record Length      137
Block Size         6144
```

See Chapter 11, "Network File System operation," on page 157 and "Log=msglevel operand" on page 171 to find out how to set the level of messages to be stored in this data set. The default message level is "info" which means that all error, attention, and informational messages are collected and stored in the z/OS NFS server log data set. Reading the message(s) in the z/OS NFS server log data set can help you identify the user correctable error or the problem error which you can report to IBM. See "Server messages" on page 241 for an explanation of the format of the messages that appear in the z/OS NFS server log data sets.

© Copyright IBM Corp. 1991, 2006

This example shows how to identify a sample problem by reading the z/OS NFS server log data set. This example shows a z/OS NFS server log data set with the message level set as "info" (informational):

```
20:25:16 GFSA305I (I) GFSAMAIN ACOPR 01 OPRPARSE: RECEIVED
   COMMAND: EXPORTFS , LENGTH = 9
20:25:16 GFSA865I (E) GFSAMAIN ANEXP 03 DOOPTION: EXPORTS:
   UNEXPECTED OPTION (PO)--EXPORTS FILE UNUSABLE.
20:25:16 GFSA866I (E) GFSAMAIN ANEXP 07 EXPORTAL: EXPORTS:
   DIRECTORY (TCP) WAS NOT EXPORTED.
20:25:16 GFSA925I (E) GFSAMAIN AXOPE 02 OPR_EXPO: ERROR WAS
   DETECTED IN THE EXPORTS FILE. EXPORT LIST NOT REBUILT
20:25:34 GFSA305I (I) GFSAMAIN ACOPR 01 OPRPARSE: RECEIVED
   COMMAND: FLUSHLOG , LENGTH = 9
   ⋮
```

The message GFSA865I in the example is an error message (indicated by the message level E).

```
20:25:16 GFSA865I (E) GFSAMAIN ANEXP 03 DOOPTION: EXPORTS:
   UNEXPECTED OPTION (PO)--EXPORTS FILE UNUSABLE.
```

The message text indicates that the option "PO" specified in the DIRECTORY statement in the exports data set is not a valid option. You could correct this error by changing "PO" to "RO" and restarting the server.

Client log data sets

The z/OS Network File System client records diagnostic information (messages and debug trace information) in the z/OS NFS client log data sets specified in the NFSCMSG1 and NFSCMSG2 DD statements of the client's startup procedure.

NFSCMSG1 is associated with the primary log, while NFSCMSG2 is associated with the secondary log. When the client is started, the primary log is used first. When the primary log is full, the logging is automatically switched or toggled to the secondary log and a console message is displayed. The last log buffer is lost when the switch takes place. When the secondary log is also full, the logging is switched back to the primary log and the original primary log content is overwritten.

The number of log records depends on the log level setting and the number of transactions that the client handles. Adjust the space allocation of your z/OS NFS client log data sets according to your installation experience to avoid frequent log switching.

Setting up the z/OS NFS client log data sets

Here is an example of setting up the z/OS NFS client log data sets:

```
//NFSCMSG1  DD   DISP=SHR,DSN=MVSNFS.CLIENT.LOG1
//NFSCMSG2  DD   DISP=SHR,DSN=MVSNFS.CLIENT.LOG2
```

File attributes of the z/OS NFS client log data sets can be set up like this:

```
Organization       PS
Record Format      VB
Record Length      137
Block Size         6144
```

Reading the message(s) in the z/OS NFS client log data set can help you identify the user correctable error or the problem error which you can report to IBM. See "Client messages" on page 281 for an explanation of the format of the messages that appear in the z/OS NFS client log data sets.

This example shows how to identify a sample problem by reading the z/OS NFS client log data set:

```
13:00:19 GFSC098I (D) CPARS 03 P_MOUNT :
   POS=32 START(7F602A78)
13:00:19 GFSC098I (D) CPARS 02 P_PARM  :
   TOKEN.POS(36)TYPE(80)VALUE(0) START(7F602A98)CUR(7F602A9D)
13:00:19 GFSC313I (I) CPARS 04 P_CHECK :
   RETRANS OPTION WILL BE IGNORED AS HARD OPTION IS ON.
13:00:19 GFSC098I (D) CVMNT 74 MI_BLDRN:
   RMN(7F2BAE30:7F172530:7F1728AC) CNT=7
   ⋮
```

The message GFSC313I in the example is an informational message (indicated by the message level I).

```
13:00:19 GFSC313I (I) CPARS 04 P_CHECK :
   RETRANS OPTION WILL BE IGNORED AS HARD OPTION IS ON.
```

The message text indicates that the keyword, *retrans*, will be ignored because *hard* has been set on. The system programmer should check the mount parameters to make sure that *hard* should be on.

Debug trace data capture

The z/OS Network File System uses trace facilities to record debug trace diagnostic information when a problem requires additional diagnostic information beyond diagnostic messages. Assuming that debug trace diagnostics were not activated at the time of the original failure, the problem must be recreated a second time and the NFS debug trace facilities turned on to capture the diagnostic information. The server provides debug trace diagnostic information from two separate trace facilities. The z/OS Network File System client provides debug trace diagnostic information from one trace facility.

z/OS NFS server debug trace capture

The z/OS NFS server records z/OS NFS debug trace diagnostic information to the z/OS component trace buffer or to the server log data sets. z/OS NFS error and informational messages are also recorded to the z/OS component trace buffer or to the server log data sets. When recording debug trace diagnostic information, IBM recommends that the z/OS component trace buffer be used to minimize the performance impact of collecting these extra diagnostics. To record z/OS NFS server debug trace diagnostic information in the z/OS component trace buffer, see "Using z/OS component tracing" on page 229.

To record z/OS NFS server debug trace diagnostic information in the log data sets, the MODIFY Operator command with the DEBUGn trace option is used.

To start NFS Server debug trace recording from the console, enter:

```
MODIFY mvsnfs,LOG=DEBUGn
```

where *mvsnfs* is the name of the z/OS NFS procedure in the PROCLIB and *n* is 1-9.

The trace records are buffered in storage. The current buffers must be flushed to disk and then the log data sets must be swapped. After swapping, the desired log data set must be copied to disk before it is overwritten by toggling back to the previous log data set by the automatic toggle mechanism.

To flush the remaining z/OS NFS buffers to disk, enter the following console command:

```
MODIFY mvsnfs,FLUSHLOG
```

where *mvsnfs* is the name of the z/OS NFS procedure in the PROCLIB.

This command enables a TSO/E user to browse all the log records that have been written by the Network File System.

To swap z/OS NFS server debug trace recording to the next log data set from the USS command line, issue:

```
MODIFY mvsnfs,SWITCHLOG
```

where *mvsnfs* is the name of the z/OS NFS procedure in the PROCLIB.

To end z/OS NFS Server debug trace recording, enter the following console command:

```
MODIFY mvsnfs,LOG=INFO
```

where *mvsnfs* is the name of the z/OS NFS procedure in the PROCLIB.

To clear (reset) the server log data sets, enter the following console command:

```
MODIFY mvsnfs,INITLOG
```

where *mvsnfs* is the name of the NFS procedure in the PROCLIB.

z/OS NFS server DEBUG trace types

The z/OS NFS server has various trace types which control the amount of trace diagnostic information recorded.

ERROR
Only writes error messages to the log data set.

INFO Writes both error and informational messages to the log data set. Some informational messages are also written to the console. See Chapter 19, "Network File System messages," on page 241for details.

DEBUGn
In addition to Error and Informational messages, some NFS activity trace diagnostic information is also recorded to the server log data set. The specific trace diagnostic information recorded depends on the debug level, specified by ***n***:

DEBUG1
Request status information is recorded: requests issued and their associated replies. Also, some locking trace diagnostic information is recorded.

DEBUG2
DEBUG1 diagnostic information, plus Network related tracing information.

DEBUG3
DEBUG2 diagnostic information, plus some key NFS trace points.

DEBUG4
DEBUG3 diagnostic information, plus detailed NFS flow information for everything, except for the HFS processing and data management functions.

DEBUG5
The Logical Layer of Data Management only.

DEBUG6
The Physical Layer of Data Management only.

DEBUG7
DEBUG1, DEBUG5, and DEBUG6 diagnostic information.

DEBUG8
The HFS processing function only. Note: This is only the NFS side of the HFS processing. It does not include any trace diagnostic information from the HFS Physical File System itself.

DEBUG9
DEBUG1 through DEBUG8 diagnostic information.

z/OS NFS client debug trace capture

The z/OS NFS client records debug trace diagnostic information and error and informational message diagnostic information to the client log data sets.

z/OS client debug trace recording can only be turned on or off; there are no graduated debug levels. Debug trace recording does not have any console command control capability. Debug trace recording is controlled from the UNIX command line. Superuser authority is required to start debug trace recording to the client log data sets.

To start z/OS NFS client debug trace recording from the USS command line, issue:

```
/usr/lpp/NFS/nfsstat -d 1
```

The current buffers must be flushed to disk and then the log data sets must be swapped. After swapping, the desired log data set must be copied to disk before it is overwritten by toggling back to the previous log data set by the automatic toggle mechanism.

To flush the remaining z/OS NFS buffer to disk, enter:

```
 /usr/lpp/NFS/nfsstat -d f
```

To swap z/OS NFS client debug trace recording to the next log data set from the USS command line, issue:

```
/usr/lpp/NFS/nfsstat -d s
```

To end z/OS NFS client debug trace recording from the USS command line, issue:

```
/usr/lpp/NFS/nfsstat -d 0
```

To clear (reset) the client log data sets from the USS command line, issue:

```
/usr/lpp/NFS/nfsstat -d c
```

Related component trace capture

The z/OS Network File System trace diagnostic information is used in parallel to other z/OS component trace diagnostic information to better understand the system activity for a problem encountered. Depending on the problem, trace diagnostic information may be needed from the z/OS components z/OS UNIX, HFS, or TCP/IP, and for an AIX client or SUN client. In addition, a z/OS dump may be required.

z/OS UNIX System Services activity trace

z/OS UNIX System Services (z/OS UNIX) diagnostic information is normally maintained in a set of cyclical buffers in memory. However, it can also be recorded to a file with a component trace. When it is maintained in memory, a dump must be taken to capture the diagnostic information for debugging purposes.

- To start the z/OS UNIX CTRACE from the console, enter:

  ```
  TRACE CT,4M,COMP=SYSOMVS
  R nn,OPTIONS=(ALL),JOBNAME=(NFS_server),END
  ```

  Note: If JOBNAME keyword is used on the statement above, be aware that z/OS UNIX uses that parm for USERID, not for job name. Specify the userid, not the jobname, for this parameter.
- Reproduce the problem in a system with minimal activity.
- Collect dump via console dump (see "z/OS dump" on page 386)
- To turn z/OS UNIX CTRACE off, issue the following console command:

  ```
  TRACE CT,OFF,COMP=SYSOMVS
  ```

z/OS Hierarchical File System (HFS) Physical File System activity trace

HFS physical file system diagnostic information is normally maintained in a set of cyclical buffers in memory. However, it can also be recorded to a file with a component trace. When it is maintained in memory, a dump must be taken to capture the diagnostic information for debugging purposes.

- To start the HFS CTRACE from the console, enter:

  ```
  TRACE CT,8M,COMP=SYSSMS
  R nn,JOBNAME=(OMVS),OPTIONS=(ENTRY,EXIT,EXITA,SPECIAL,CB,RRTN,COMP=(PFS)),END
  ```
- Reproduce the problem in a system with minimal activity.
- Collect dump via console dump (see"z/OS dump" on page 386)
- To turn HFS CTRACE off, issue the following console command:

  ```
  TRACE CT,OFF,COMP=SYSSMS
  ```

z/OS TCP/IP activity trace

The TCP/IP CTRACE and Packet Trace are normally maintained in a set of cyclical buffers in memory. However, they can also be recorded to a file with a component trace. When they is maintained in memory, a dump must be taken to capture the diagnostic information for debugging purposes.

- To start the TCP/IP CTRACE from the console, enter:

  ```
  TRACE CT,ON,COMP=SYSTCPIP,SUB=(tcpipprocname)
  R XX,JOBNAME=(nfsprocname,tcpipprocname)
  R XX,OPTIONS=(ENGINE,PFS,SOCKET,INTERNET,TCP,UDP,IOCTL),END
  ```
- To start the TCP/IP packet trace from the console, enter:

```
TRACE CT,ON,COMP=SYSTCPDA,SUB=(tcpipprocname)
V TCPIP,tcpipprocname,PKTTRACE,ON,ABBREV=152,IP=xx.xx.xx.xx
```

Note: xx.xx.xx.xx is the client_IP_address.

- Reproduce the problem in a system with minimal activity.
- To collect a dump of TCPIP Dataspace and TCPIP address space, issue the following console command:

```
DUMP COMM=('text')
R xx,JOBNAME=(tcpipprocname,nfsprocname),DSPNAME=('tcpipprocname'.*),
SDATA=(ALLNUC,PSA,GRSQ,SUM,CSA,LPA,LSQA,RGN,SWA,SQA,TRT),END
```

NOTE: See "z/OS dump" on page 386)

- To turn TCP/IP CTRACE off, issue the following console command:

```
TRACE CT,OFF,COMP=SYSTCPIP,SUB=(tcpipprocname)
V TCPIP,tcpipprocname,PKT,OFF,IP=xx.xx.xx.xx
```

- To stop the TCPIP Packet trace, issue the following console command:

```
TRACE CT,OFF,COMP=SYSTCPDA,SUB=(tcpipprocname)
```

NOTE: The first step to collecting traces to the dataspace is to ensure that the *bufsize* in SYS1.PARMLIB member CTIEZB00 is set to at least 8 MB. It may need to be set higher depending on the amount of trace diagnostic information desired, but 8 MB should be used as a minimum. Using the maximum buffer size of 256 MB, you will be sure that the maximum amount of contiguous trace diagnostic information is captured. A 256 MB buffer size is recommended.

TCP/IP will need to be restarted for the change in buffer size to take affect. The trace diagnostic information will be captured using an z/OS dump command that will dump the TCP/IP dataspace named TCPIPDS1. Be aware that this method may result in lost trace diagnostic information as the possibility of wrapping is high. The dump command should be issued soon after the problem occurs.

AIX client activity trace

The AIX IP trace contains a trace of the IP activity from the AIX client:

```
rm /tmp/ibmsupt
```

NOTE: Make sure that there is enough space under /tmp for the IP-traces.

Systems with high TCP/IP activity may use up to 20 MB per minute.

snap -gkitn (takes 2-3 minutes and collects all necessary AIX diagnostic information).

- To start the AIX IP trace, issue the following AIX command:

```
startsrc -s iptrace -a "-s -d -b
/tmp/ibmsupt/testcase/iptrc.bin"

cd /tmp/ibmsupt/testcase
script cmd_log (puts all keybaord actions to file cmd_log in current directory)
Ping -c 5 (z/OS NFS server)  showmount -e
cd to NFS mounted z/OS filessysem
cp /etc/filesystems /tmp/ibmsupt/testcase/filesystems.out
mount
```

- Reproduce the problem in a system with minimal activity.
- To stop the AIX IP Trace, issue the following AIX command:

```
stopsrc -s iptrace
```

- To format the Trace Report, issue the following AIX command:

```
ipreport -rns /tmp/ibmsupt/testcase/iptrc.bin > /tmp/ibmsupt/testcase/iptrc.out
```

NOTE: Formatting of the IP trace takes about 5 times.

- To stop script recording, issue the following console command:

  ```
  CTRL-D
  ```

- To generate snap.pax.Z file in /tmp/ibmsupt directory, issue the following console command:

  ```
  snap -c
  ```

- Send snap.pax.Z file to NFS service personnel.

SUN client activity trace

The SUN snoop trace contains a trace of the IP activity from the SUN client.

- To start the SUN snoop trace, issue the following SUN command:

  ```
  snoop -o my.trace myclient myserver
  ```

- To stop the SUN snoop Trace, issue the following SUN command:

  ```
  CTRL-C
  ```

- To format Trace Report, issue the following SUN command:

  ```
  snoop -i my.trace -v >my.trace.report
  ```

z/OS dump

A z/OS memory dump contains the current state of the machine. The actual contents of the dump will depend on the address spaces and data spaces selected. The dump will also include any component traces running at the time the dump is taken. IBM generally recommends taking a synchronous dump instead of an asynchronous dump. A synchronous dump ensures that the memory contents do not change between the time the dump request is made and the time the dump is actually taken on the machine.

To collect a console dump of NFS server, NFS client, z/OS UNIX, and TCPIP, issue the following console command:

```
DUMP COMM=(description)

R nn,JOBNAME=(OMVS,TCPproc,NFSSproc,NFSCproc,jobname*),CONT

R nn,DSPNAME=('OMVS'.*,'tcpproc'.*),CONT

R nn, SDATA=(PSA,SQA,LSQA,RGN,TRT,LPA,CSA,GRSQ,SUM,ALLNUC),CONT
```

Appendix M. GFSAMHDJ sample code for creating NFS mount handle data sets and lock data sets

The following sample code (GFSAMHDJ member of SYS1.SAMLIB) shows how to allocate mount handle data sets, copy previous mount handle data sets to the new size needed for z/OS V1R7, and how to allocate lock data sets.

```
//MVSNFSA  JOB  ,
//    MSGCLASS=A,MSGLEVEL=(1,1),TIME=30,REGION=12M,CLASS=A
//*
//*  FOR MOUNT HANDLE DATABASE DEFINITION, USE:
//*    KEYS(16 0) -
//*  FOR LOCKING DATABASE DEFINITION, USE
//*    KEYS(8 0) -
//*
//*************************************************************
//*
//*
//*  CREATE A VSAM KSDS DATASET FOR THE MOUNT HANDLE DATABASE
//*
//*  THIS STEP IS ONLY REQUIRED IF THE OLD MOUNT HANDLE DATASETS
//*  DO NOT HAVE THE SAME RECORD SIZE DEFINITIONS AS THE CURRENT
//*  DEFINITIONS OR NO MOUNT HANDLE DATASETS EXIST.
//*
//*  REPLACE THE FOLLOWING FIELDS BELOW:
//*
//*     MVSNFS.FHDBASE   WITH DESIRE MOUNT HANDLE DATASET 1 NAME
//*     xxxxxx           WITH VOLUME SERIAL FOR ALLOCATION
//*
//*
//*************************************************************
//DEFINE1 EXEC PGM=IDCAMS,REGION=512K,COND=(EVEN)
//SYSPRINT DD  SYSOUT=*
//SYSIN DD *
 DEFINE CLUSTER ( NAME(MVSNFS.FHDBASE) -
             VOL(xxxxxx) -
             CYL (1 1) -
             INDEXED -
             REUSE -
             KEYS(16 0) -
             SHAREOPTIONS(1 3) -
             RECSZ(1700 2000) )
 LISTC ENT(MVSNFS.FHDBASE) ALL
/*
```

Figure 42. Sample code for creating mount handle data sets and lock data sets (Part 1 of 3)

© Copyright IBM Corp. 1991, 2006

```
//*************************************************************
//*
//*
//*  CREATE THE SECOND VSAM KSDS FOR THE MOUNT HANDLE DATABASE
//*         ON A DIFFERENT VOLUME
//*
//*  THIS STEP IS ONLY REQUIRED IF THE OLD MOUNT HANDLE DATASETS
//*  DO NOT HAVE THE SAME RECORD SIZE DEFINITIONS AS THE CURRENT
//*  DEFINITIONS OR NO MOUNT HANDLE DATASETS EXIST.
//*
//*  REPLACE THE FOLLOWING FIELDS BELOW:
//*
//*     MVSNFS.FHDBASE2  WITH DESIRE MOUNT HANDLE DATASET 2 NAME
//*     yyyyyy           WITH VOLUME SERIAL FOR ALLOCATION
//*
//*
//*************************************************************
//DEFINE2 EXEC PGM=IDCAMS,REGION=512K,COND=(EVEN)
//SYSPRINT DD  SYSOUT=*
//SYSIN DD *
 DEFINE CLUSTER ( NAME(MVSNFS.FHDBASE2) -
             VOL(yyyyyy) -
             CYL (1 1) -
             INDEXED -
             REUSE -
             KEYS(16 0) -
             SHAREOPTIONS(1 3) -
             RECSZ(1700 2000) )
 LISTC ENT(MVSNFS.FHDBASE2) ALL
/*
//*************************************************************
//*
//*                                                        @01A
//*  COPY OLD SMALLER RECORD SIZE MOUNT HANDLE DATABASES TO THE
//*       NEW LARGER  RECORD SIZE MOUNT HANDLE DATABASES
//*       REQUIRED AS OF Z/OS NFS 1.7
//*
//*  THIS STEP IS ONLY REQUIRED IF THE OLD MOUNT HANDLE DATASETS
//*  DO NOT HAVE THE SAME RECORD SIZE DEFINITIONS AS THE CURRENT
//*  DEFINITIONS.
//*
//*
//*  REPLACE THE FOLLOWING FIELDS BELOW:
//*
//*     MVSNFS1.FHDBASE  WITH OLD MOUNT HANDLE DATASET 1 NAME
//*     MVSNFS1.FHDBASE2 WITH OLD MOUNT HANDLE DATASET 2 NAME
//*     MVSNFS.FHDBASE   WITH NEW MOUNT HANDLE DATASET 1 NAME
//*     MVSNFS.FHDBASE2  WITH NEW MOUNT HANDLE DATASET 2 NAME
//*
//*
//*************************************************************
//IDCAMS1 EXEC PGM=IDCAMS,COND=(EVEN)
//SYSPRINT  DD SYSOUT=*
//OLDFHDB1  DD DSN=MVSNFS1.FHDBASE,DISP=SHR
//OLDFHDB2  DD DSN=MVSNFS1.FHDBASE,DISP=SHR
//NEWFHDB1  DD DSN=MVSNFS.FHDBASE,DISP=SHR
//NEWFHDB2  DD DSN=MVSNFS.FHDBASE,DISP=SHR
//SYSIN     DD *
 REPRO INFILE(OLDFHDB1) OUTFILE(NEWFHDB1) REUSE
 REPRO INFILE(OLDFHDB2) OUTFILE(NEWFHDB2) REUSE
/*
```

Figure 42. Sample code for creating mount handle data sets and lock data sets (Part 2 of 3)

```
//*************************************************************
//*
//*                                                        @01A
//*  CREATE A VSAM KSDS DATASET FOR THE LOCKING DATABASE
//*
//*  THE LOCK DATA SETS MUST ALWAYS BE ALLOCATED EVEN IF NONLM
//*  IS SPECIFIED IN THE SITE ATTRIBUTES.
//*
//*  REPLACE THE FOLLOWING FIELDS BELOW:
//*
//*     MVSNFS.LDBASE    WITH DESIRE LOCK DATASET 1 NAME
//*     xxxxxx           WITH VOLUME SERIAL FOR ALLOCATION
//*
//*
//*************************************************************
//DEFINE3 EXEC PGM=IDCAMS,REGION=512K,COND=(EVEN)
//SYSPRINT DD  SYSOUT=*
//SYSIN DD *
 DEFINE CLUSTER ( NAME(MVSNFS.LDBASE) -
             VOL(xxxxxx) -
             CYL (1 1) -
             INDEXED -
             REUSE -
             KEYS(8 0) -
             SHAREOPTIONS(1 3) -
             RECSZ(1700 2000) )
 LISTC ENT(MVSNFS.LDBASE) ALL
/*
//*************************************************************
//*
//*                                                        @01A
//*  CREATE THE SECOND VSAM KSDS FOR THE LOCKING DATABASE
//*         ON A DIFFERENT VOLUME
//*
//*  THE LOCK DATA SETS MUST ALWAYS BE ALLOCATED EVEN IF NONLM
//*  IS SPECIFIED IN THE SITE ATTRIBUTES.
//*
//*  REPLACE THE FOLLOWING FIELDS BELOW:
//*
//*     MVSNFS.LDBASE2   WITH DESIRE LOCK DATASET 2 NAME
//*     yyyyyy           WITH VOLUME SERIAL FOR ALLOCATION
//*
//*
//*************************************************************
//DEFINE4 EXEC PGM=IDCAMS,REGION=512K
//SYSPRINT DD  SYSOUT=*
//SYSIN DD *
 DEFINE CLUSTER ( NAME(MVSNFS.LDBASE2) -
             VOL(yyyyyy) -
             CYL (1 1) -
             INDEXED -
             REUSE -
             KEYS(8 0) -
             SHAREOPTIONS(1 3) -
             RECSZ(1700 2000) )
 LISTC ENT(MVSNFS.LDBASE2) ALL
/*
//*
```

Figure 42. Sample code for creating mount handle data sets and lock data sets (Part 3 of 3)

Appendix N. Setting up a Kerberos Key Distribution Center

In order to start a z/OS NFS server with Kerberos authentication features, a Kerberos Key Distribution Center must be ready before the z/OS NFS server starts up. This section lists the basic steps involved in setting up the z/OS KDC which will be compatible with the z/OS NFS server environment. For more advanced configurations and detailed explanations of the setup steps and the reasoning behind, please refer to *z/OS Integrated Security Services Network Authentication Service Administration*. For other platform's KDC setups, please consult the specific platform documentation.

These steps assume that Resource Access Control Facility (RACF) is available in the system. If you have a different but equivalent external security manager, please refer to the documentation of the product for instructions.

1. Copy the SKRBKDC started task procedure from EUVF.SEUVFSAM to SYS1.PROCLIB. The SYS1.PROCLIB(SKRBKDC) should look like the following:

```
//*********************************************************************
//*
//* Procedure for starting the Kerberos SKRBKDC started task
//* Specify PARMS='-kdc' to enable the Kerberos KDC services.
//* Specify PARMS='-nokdc' to disable the Kerberos KDC services.
//*
//*********************************************************************
//SKRBKDC  PROC  REGSIZE=256M,OUTCLASS='A',PARMS='-kdc'
//*--------------------------------------------------------------------
//GO       EXEC  PGM=EUVFSKDC,REGION=&REGSIZE,TIME=1440,
//  PARM=('ENVAR("LANG=En_US.IBM-1047"),TERM(DUMP) / &PARMS      X
//             1>DD:STDOUT 2>DD:STDERR')
//STDOUT   DD  SYSOUT=&OUTCLASS,DCB=LRECL=250,
//  FREE=END,SPIN=UNALLOC
//STDERR   DD  SYSOUT=&OUTCLASS,DCB=LRECL=250,
//  FREE=END,SPIN=UNALLOC
//SYSOUT   DD  SYSOUT=&OUTCLASS,
//  FREE=END,SPIN=UNALLOC
//CEEDUMP  DD  SYSOUT=&OUTCLASS,
//  FREE=END,SPIN=UNALLOC
```

2. Copy the sample Kerberos configuration file in z/OS UNIX from /usr/lpp/skrb/examples/krb5.conf to /etc/skrb/krb5.conf. The permission bits of this file should allow only the administrator to modify it but everyone else to be able to read.

   Note: Currently Linux (for example, Enterprise Linux 4) does not support the des-cbc-md5 encryption type. If the z/OS NFS server will be supporting multiple platforms of NFS client, IBM recommends using des-cbc-crc encryption types only, as shown in the copy below.

```
;--------------------------------------------------------------------;
;  Sample Kerberos configuration file                                ;
;                                                                    ;
;  Copy this file to /etc/skrb/krb5.conf and then tailor it for      ;
;  your Kerberos configuration                                       ;
;                                                                    ;
;  Do not enable DES3 encryption unless all of the systems in the    ;
;  realm have DES3 support.  In order to use DES3 encryption for     ;
;  tickets, you must set the SKDC_TKT_ENCTYPES environment variable  ;
;  in /etc/skrb/home/kdc/envar.                                      ;
;--------------------------------------------------------------------;
 [libdefaults]
```

© Copyright IBM Corp. 1991, 2006

```
default_realm = KRB390.IBM.COM
kdc_default_options = 0x40000010
use_dns_lookup = 0
; Default encryption types if DES3 is not supported
default_tkt_enctypes = des-cbc-crc
default_tgs_enctypes = des-cbc-crc

[realms]
KRB390.IBM.COM = {
kdc = dcesec4.krb390.ibm.com:88
kpasswd_server = dcesec4.krb390.ibm.com:464
admin_server = dcesec4.krb390.ibm.com:749
}

KRB2000.IBM.COM = {
kdc = sstone1.krb2000.ibm.com:88
admin_server = sstone1.krb2000.ibm.com:749
}

[domain_realm]
.krb390.ibm.com = KRB390.IBM.COM
.krb2000.ibm.com = KRB2000.IBM.COM

[capaths]
KRB390.IBM.COM = {
KRB2000.IBM.COM = .
}
```

3. Copy the environment variable definitions from /usr/lpp/skrb/examples/skrbkdc.envar to /etc/skrb/home/kdc/envar. Depending on which type of KDC is being set up, the environment variable SKDC_DATABASE should be set to SAF or NDBM accordingly (default is set to SAF registry type KDC). The file permissions should allow only the administrator to read and update.
4. Add the path "PATH=/usr/lpp/skrb/bin:$PATH" in the z/OS UNIX to the user's ".profile" and export the "PATH".
5. Issue the following Ishell commands (entering each command on a single line). If needed, consult *z/OS Integrated Security Services Network Authentication Service Administration* for explanations.
   - `RDEFINE FACILITY IRR.RUSERMAP UACC(read)`
   - `SETROPTS RACLIST(facility) REFRESH`
   - `AU SKRBKDC DFLTGRP(sys1) NOPASSWORD OMVS(uid(0) PROGRAM('/bin/sh') HOME('/etc/skrb/home/kdc'))`
   - `RDEFINE REALM KERBDFLT KERB(kerbname(REALM) PASSWORD(password) MINTKTLFE(15) DEFTKTLFE(36000) MAXTKTLFE(86400))`
6. For SAF registry KDC, continue with the following steps. For NDBM registry KDC, skip to step 7 on page 393.
   a. Issue the following Ishell commands, entering each command on a single line:
      - `SETROPTS CLASSACT(appl) RACLIST(appl)`
      - `RDEFINE APPL SKRBKDC UACC(read)`
      - `SETROPTS CLASSACT(ptktdata) RACLIST(ptktdata)`
      - `RDEFINE PTKTDATA SKRBKDC UACC(none) SSIGNON(keymasked(3734343237343131))`
      - `SETROPTS RACLIST(appl ptktdata) REFRESH`
      - `SETROPTS GENERIC(started)`
      - `RDEFINE STARTED SKRBKDC.** STDATA(user(skrbkdc))`
      - `RDEFINE STARTED SKRBWTR.** STDATA(user(skrbkdc))`
      - `SETROPTS RACLIST(started) REFRESH`
      - `AU KADMIN DFLTGRP(sys1) PASSWORD(password) KERB(kerbname(kadmin/admin))`
      - `ALU KADMIN PASSWORD(password) NOEXPIRED`
      - `AU CHANGEPW DFLTGRP(sys1) PASSWORD(password) KERB(kerbname(kadmin/changepw))`
      - `ALU CHANGEPW PASSWORD(password) NOEXPIRED`

b. Create NFS principal for the z/OS NFS server.

```
AU racfid PASSWORD(password)  KERB(KERBNAME(nfs/host.domain))

ALU racfid PASSWORD(password) NOEXPIRED
```

Note: The ALTUSER command converts the password to upper case if the MIXEDCASE SETROPTS option is not set. If MIXEDCASE is not set, you must ensure that the uppercase value is used when you request an initial ticket. The principal name is not converted to upper case and the realm name is not included. You must change the password for the user in order to create the Kerberos secret key.

c. To create keytab for z/OS NFS server, issue, for example, the following command in z/OS UNIX.

```
keytab add -p password -k /path/temp.keytab -v 1 nfs/host.domain
```

Note: The version value (-v option) and password need to match the version and password used when creating the NFS principal.

d. For multiple Kerberos realms environment, create ticket-granting tickets in Ishell:

```
RDEFINE REALM /.../REALM1/KRBTGT/REALM2/  KERB(PASSWORD(password))

RDEFINE REALM /.../REALM2/KRBTGT/REALM1/  KERB(PASSWORD(password))
```

Note: the /.../ and trailing slash are required.

e. Add Kerberos segments to existing user definitions. These Kerberos segments serve as the Kerberos principals in the Kerberos database.

To add a RACF id, issue Ishell command, for example:

```
AU (RACFID1) OWNER (IBMUSER) OMVS(UID(101))
```

To define Kerberos segment to the above user definition, issue Ishell command:

```
ALTUSER RACFID1 PASSWORD(password) NOEXPIRED    KERB(KERBNAME(user1))
```

f. Start the skrbkdc task.

g. Continue to step 8 on page 394 to complete KDC setup.

7. For NDBM registry type KDC, follow these steps.

a. To create initial registry database files, issue z/OS UNIX command:

```
kdb5_ndbm create
```

IBMUSER and IBMUSER/admin user principals are now created with initial password of IBMUSER.

b. Copy sample KDC configuration file from /usr/lpp/skrb/examples/kdc.conf to /etc/skrb/home/kdc/kdc.conf and set the values inside as needed or leave them to default values.

c. Copy the sample administration access control file from /usr/lpp/skrb/examples/kadm5.acl to /etc/skrb/home/kdc/kadm5.acl . The administrator can choose to customize it or leave it as default.

d. Start the skrbkdc task.

e. Create NFS principal for the z/OS NFS server using the kadmin interface in z/OS UNIX.

To enter the kadmin interface, issue z/OS UNIX command:

```
kadmin -p IBMUSER/admin -w IBMUSER
```

To create NFS principle, enter:

```
kadmin> addprinc nfs/host.domain
```

f. For multiple Kerberos realm environment, create ticket-granting tickets in kadmin interface:

```
kadmin> addprinc -e des-cbc-crc:normal  krbtgt/REALM1@REALM2

kadmin> addprinc -e des-cbc-crc:normal  krbtgt/REALM2@REALM1
```

Note: The "-e des-cbc-crc:normal" should be used in order to support various UNIX platforms. The passwords specified for these 2 principals should be the same.

g. Create keytab for the z/OS NFS server by issuing in kadmin interface:

```
kadmin> ktadd -e des-cbc-crc:normal -k /path/temp.keytab     nfs/host.domain
```

Note: The "-e des-cbc-crc:normal" should be used in order to support various UNIX platforms. The passwords specified for these 2 principals should be the same.

h. Add Kerberos principals into the Kerberos database through the kadmin interface:

```
kadmin> addprinc user1
kadmin> addprinc user2
```

8. Put (or ftp in binary mode if different hosts) the keytab on z/OS NFS server's /etc/skrb directory and rename it to /etc/skrb/krb5.keytab, in order to be able to start the z/OS NFS server later.

Appendix O. Accessibility

Accessibility features help a user who has a physical disability, such as restricted mobility or limited vision, to use software products successfully. The major accessibility features in z/OS enable users to:

- Use assistive technologies such as screen readers and screen magnifier software
- Operate specific or equivalent features using only the keyboard
- Customize display attributes such as color, contrast, and font size

Using assistive technologies

Assistive technology products, such as screen readers, function with the user interfaces found in z/OS. Consult the assistive technology documentation for specific information when using such products to access z/OS interfaces.

Keyboard navigation of the user interface

Users can access z/OS user interfaces using TSO/E or ISPF. Refer to *z/OS TSO/E Primer*, *z/OS TSO/E User's Guide*, and *z/OS ISPF User's Guide Vol I*for information about accessing TSO/E and ISPF interfaces. These guides describe how to use TSO/E and ISPF, including the use of keyboard shortcuts or function keys (PF keys). Each guide includes the default settings for the PF keys and explains how to modify their functions.

z/OS information

z/OS information is accessible using screen readers with the BookServer/Library Server versions of z/OS books in the Internet library at:

`www.ibm.com/servers/eserver/zseries/zos/bkserv/`

© Copyright IBM Corp. 1991, 2006

Notices

This information was developed for products and services offered in the U.S.A. IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described 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:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:
INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

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:

© Copyright IBM Corp. 1991, 2006

IBM Corporation
Mail Station P300
2455 South Road
Poughkeepsie, NY 12601-5400
USA

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this information and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement, or any equivalent agreement between us.

Programming interface information

This document documents intended Programming Interfaces that allow the customer to write programs to obtain the services of z/OS NFS.

Trademarks

The following terms are trademarks of the IBM Corporation in the United States, or other countries, or both:

IBM
AIX
BookManager
DFSMS/MVS
DFSMSdfp
DFSMSdss
DFSMShsm
IBMLink
Language Environment
MVS
MVS/DFP
MVS/ESA
OS/2
OS/390
RACF
RETAIN
Resource Link
RISC System/6000
RS/6000
VTAM
z/OS
z/VM
zSeries

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product, and service names, which may be trademarks or service marks of others.

Glossary

This glossary defines technical terms and abbreviations used in DFSMS documentation. If you do not find the term you are looking for, refer to the index of the appropriate DFSMS manual or view the *Glossary of Computing Terms* located at:

http://www.ibm.com/ibm/terminology/

This glossary includes terms and definitions from:

- The *American National Standard Dictionary for Information Systems*, ANSI X3.172-1990, copyright 1990 by the American National Standards Institute (ANSI). Copies may be purchased from the American National Standards Institute, 11 West 42nd Street, New York, New York 10036. Definitions are identified by the symbol (A) after the definition.
- The *Information Technology Vocabulary* developed by Subcommittee 1, Joint Technical Committee 1, of the International Organization for Standardization and the International Electrotechnical Commission (ISO/IEC JTC1/SC1). Definitions of published part of this vocabulary are identified by the symbol (I) after the definition; definitions taken from draft international standards, committee drafts, and working papers being developed by ISO/IEC JTC1/SC1 are identified by the symbol (T) after the definition, indicating that final agreement has not yet been reached among the participating National Bodies of SC1.
- The *IBM Dictionary of Computing*, New York: McGraw-Hill, 1994.

The following cross-reference is used in this glossary:

See: This refers the reader to (a) a related term, (b) a term that is the expanded form of an abbreviation or acronym, or (c) a synonym or more preferred term.

A

access. To obtain computing services.

access method. (1) A mainframe data management routine that moves data between storage and an I/O device in response to requests made by a program. (2) The part of the distributed data management architecture which accepts commands to access and process the records of a file.

access permission. A group of designations that determine who can access a particular AIX or UNIX file and how the user can access the file.

ACS. See *automatic class selection*.

ACS routine. See *automatic class selection (ACS) routine*.

address. The unique identifier assigned to each device or workstation connected to a network.

address space. The complete range of addresses in memory available to a computer program.

Advanced Interactive Executive (AIX). IBM's licensed version of the UNIX operating system.

AIX. See *Advanced Interactive Executive*.

alias. An alternative name for a catalog, a non-VSAM data set, or a member of a partitioned data set (PDS) or partitioned data set extended (PDSE).

alias entry. An entry that relates an alias to the real entry name of a user catalog or non-VSAM data set.

allocation. (1) Generically, the entire process of obtaining a volume and unit of external storage, and setting aside space on that storage for a data set. (2) The process of connecting a program to a data set or devices.

American Standard Code for Information Interchange (ASCII). The standard code used for information exchange among data processing systems, data communication systems, and associated equipment. It uses a coded character set consisting of 7-bit coded characters.

APAR. See *authorized program analysis report*.

APF. Authorized program facility.

API. See *application programming interface*.

application programming interface (API). A formally defined programming language interface between an IBM system control program or a licensed program and the user of a program.

ASCII. See *American National Standard Code for Information Interchange*.

atime. The time when the file was last accessed.

© Copyright IBM Corp. 1991, 2006

automatic class selection (ACS). A mechanism for assigning Storage Management Subsystem classes and storage groups to data sets.

automatic class selection (ACS) routine. A procedural set of ACS language statements. Based on a set of input variables, the ACS language statements generate the name of a predefined SMS class, or a list of names of predefined storage groups, for a data set.

authorized program analysis report (APAR). A request for correction of a problem caused by a suspected defect in a current unaltered release of a program.

B

basic sequential access method (BSAM). An access method for storing or retrieving data blocks in a continuous sequence, using either a sequential access or a direct access device.

BIOD. The caching daemon that caches directory lookups and file data when remote files are accessed from the host.

block. A string of data elements recorded, processed, or transmitted as a unit. The elements can be characters, words, or physical records.

BSAM. See *Basic sequential access method.*

C

CCSID. See *Coded Character Set Identifier*.

CDRA. Character data representation architecture.

Character Data Representation Architecture (CDRA) API. A set of identifiers, services, supporting resources, and conventions for consistent representation, processing, and interchange of character data.

client. (1) A user. (2) A consumer of resources or services. (3) A functional unit that receives shared services from a server. (4) A system that is dependent on a server to provide it with programs or access to programs. (5) On a network, the computer requesting services or data from another computer.

client-server relationship. Any process that provides resources to other processes on a network is a *server*. Any process that employs these resources is a *client*. A machine can run client and server processes at the same time.

Coded Character Set Identifier (CCSID). A 16-bit number that identifies a specific encoding scheme identifier, character set identifiers, code page identifiers, and additional coding required information. The CCSID uniquely identifies the coded graphic character representation used.

connection. An association established between functional units for conveying information.

current directory. The currently active directory; the directory that is searched when you enter a file name without indicating the directory that contains the file name. When you specify a file name without specifying a directory, the system assumes that the file is in the current directory.

D

daemon. A background process that is initiated at system start that continuously performs a function required by another process.

DA. Direct access.

DAE. Dump Analysis and Elimination.

DAM. Direct access method.

DASD. Direct access storage device.

DASD volume. A Direct Access Storage Device space identified by a common label and accessed by a set of related addresses. See also *volume*, *primary storage*, *migration level 1*, *migration level 2*.

data control block (DCB). A control block used by access method routines in storing retrieving data.

data set. In DFSMS, the major unit of data storage and retrieval, consisting of a collection of data in one of several prescribed arrangements and described by control information to which the system has access. In z/OS non-UNIX environments, the terms *data set* and *file* are generally equivalent and sometimes are used interchangeably. See also *file*. In z/OS UNIX environments, the terms *data set* and *file* have quite distinct meanings.

data set organization. The type of arrangement of data in a data set. Examples are sequential organization or partitioned organization.

DBCS. Double byte character set.

DCB. See *data control block*.

DES authentication. Requires a client to send credentials (its name, conversation key, window key, and a time stamp) to the server. The server then returns a verifier to the client. Data Encryption Standard (DES) credentials are sometimes called “secure” credentials because they are based on a sender’s ability to encrypt data using a common time reference. DES-style credentials require a randomly generated key used to encrypt a common reference time which is then used in

turn to create a *conversation key*. DES-style credentials are less vulnerable to penetration than UNIX-style credentials because of the generated conversation key and the data that is encrypted with the conversation key. Other forms of attack (for example, replaying message traffic) are resisted because of the timestamp in the transmission. The Network File System does not support DES-style credentials.

DFSMSdfp. A DFSMS functional component or base element of z/OS, that provides functions for storage management, data management, program management, device management, and distributed data access.

direct access file. A type of MVS file for storing data on a random access device that is accessed using a record address.

directory. A file that maps the names of other directories and files to their locations.

DSCB . Data set control block.

E

EBCDIC. See *extended binary-coded decimal interchange code*.

end user. A person in a data processing installation who requires the services provided by the computer system.

EOR. End of record.

entry-sequenced data set (ESDS). In VSAM, a data set whose records are loaded without respect to their contents and whose RBAs cannot change. Records are retrieved and stored by addressed access, and new records are added at the end of the data set.

ESDS. See *entry-sequenced data set (ESDS)*.

ESS. Enterprise Storage Server.

ETR. Electronic technical response.

exports data set. An MVS file on the server containing entries for directories that can be exported to Network File System clients. It is used by the server to determine which MVS files and prefixes can be mounted by a client, and to write-protect MVS files on the server.

extended binary-coded decimal interchange code (EBCDIC). A coded character set consisting of 8-bit coded characters.

External Data Representation (XDR). A standard developed by Sun Microsystems, Incorporated for representing data in machine-independent format. XDR is a vendor independent way of representing the data. By using the XDR standard data representation convention, systems do not have to understand and translate every data format that exists on the network; there is only the one convention. Data is translated into XDR format before it is sent over the network and, at the reception point, is translated into the data convention used there. This means that new computer architectures can be integrated into the network without requiring the updating of translation routines. The new architecture simply includes a routine that translates its data format into XDR format and the new member of the network is ready to go.

Using XDR, data can be accessed or exchanged among machines of various hardware and software architectures without any translation or interpretation problems. Word lengths, byte ordering, and floating point representations appear to be the same to all nodes in the network.

F

FDDI. Fiber Distributed Data Interface (FDDI).

file. A collection of information treated as a unit. In non-UNIX System Services MVS environments, the terms *data set* and *file* are generally equivalent and are sometimes used interchangeably. See also *data set*.

file handle. A file handle is used by the client and server sides of the Network File System to specify a particular file or prefix.

file system. In the z/OS UNIX HFS environment, the collection of files and file management structures on a physical or logical mass storage device, such as a diskette or minidisk. See also *HFS data set*.

File Transfer Protocol (FTP). A TCP/IP protocol used for transferring files to and from foreign hosts. FTP also provides the capability to access directories.

FMID. Function modification identifier.

FUB. File usage block.

foreign host. Any host on the network other than the local host.

free space. Space reserved within the control intervals of a key-sequenced data set for inserting new records into the data set in key sequence or for lengthening records already there; also, whole control intervals reserved in a control area for the same purpose.

FTP. File Transfer Protocol.

G

gateway. A functional unit that interconnects two computer networks with the different network architectures. A gateway connects networks or systems

of different architectures. A bridge interconnects networks or systems with the same or similar architectures.

GID. See *group number (GID)*.

group. (1) With respect to partitioned data sets, a member and the member's aliases that exist in a PDS or PDSE, or in an unloaded PDSE. (2) A collection of users who can share access authorities for protected resources.

group number (GID). A unique number assigned to a group of related users. The group number can often be substituted in commands that take a group name as an argument.

GXB . Global exit block.

H

handle. (1) In the Advanced DOS and OS/2 operating systems, a binary value created by the system that identifies a drive, directory, and file so that the file can be found and opened. (2) In the AIX operating system, the data structure that is a temporary local identifier on an object. Allocating a handle creates it. Binding a handle makes it identify an object at a specific location. (3) In the OS/400 application programming interfaces, a variable that represents an object.

HFS. See *hierarchical file system (HFS)*.

hierarchical file system (HFS) data set. A data set that contains a POSIX-compliant hierarchical file system, which is a collection of files and directories organized in a hierarchical structure, that can be accessed using the UNIX System Services MVS facilities.

host. A computer connected to a network that provides an access method to that network. A host provides end-user services.

I

ICP. Interconnect Control Program (ICP).

IDCAMS. Integrated catalog access methods services.

Internet. A specific internetwork that includes ARPANET, MILNET, and NSFnet. These networks use the TCP/IP protocol suite. See also *internetwork*.

Internet Protocol (IP). The TCP/IP layer between the higher-level host-to-host protocol and the local network protocols. IP uses local area network protocols to carry packets in the form of diagrams to the next gateway or destination host.

internetwork. A collection of packet-switched networks that are connected by gateways. They work as a single network.

interoperability. The ability of hardware and software from different vendors to communicate on a network.

inter-process communication (IPC). Ways for programs to communicate data to each other and to synchronize their activities. Semaphores, signals, and internal message queues are common methods of inter-process communication.

IP. Internet Protocol.

IPC. See *inter-process communication (IPC)*.

J

JCL. See *job control language*.

Job control language (JCL). A problem-oriented language used to identify the job or describe its requirements to an operating system.

K

key-sequenced data set (KSDS). A VSAM data set whose records are loaded in key sequence and controlled by an index. Records are retrieved and stored by keyed access or by addressed access, and new records are inserted in the data set in key sequence because of free space allocated in the data set. Relative byte addresses of records can change because of control interval or control area splits.

KSDS. See *key-sequenced data set (KSDS)*.

L

LFS. Logical file system.

local host. The computer to which a user's terminal is directly connected.

local network. That portion of a network physically connected to the host without intermediate gateways.

M

management class. A collection of management attributes, defined by the storage administrator, used to control the release of allocated but unused space; to control the retention, migration, and backup of data sets; to control the retention and backup of aggregate groups, and to control the retention, backup, and class transition of objects.

master catalog. A key-sequenced data set or file with an index containing extensive data set and volume information that VSAM requires to locate data sets or

files, to allocate and deallocate storage space, to verify the authorization of a program or operator to gain access to a data set or file, and to accumulate usage statistics for data sets or files.

MBCS. Multiple byte character set.

migration level 1. DFSMShsm-owned DASD volumes that contain data sets migrated from primary storage volumes. The data can be compressed. See also *storage hierarchy*. Contrast with *primary storage, migration level 2*.

migration level 2. DFSMShsm-owned tape or DASD volumes that contain data sets migrated from primary storage volumes or from migration level 1 volumes. The data can be compressed. See also *storage hierarchy*. Contrast with *primary storage, migration level 1*.

mkdir. An AIX, UNIX, OS/2 and DOS command used to make a new directory.

mount. A host-linked operation which results in a tape cartridge being physically inserted into a tape drive.

mount handle data set. A data set used to store the file handles of mount points.

mount point. A directory established in a workstation or a server local directory that is used during the transparent accessing of a remote file.

mtime. The time of last data modification.

MTU. Maximum transmission unit (MTU).

MVS. Multiple virtual system.

MVS/ESA. Multiple Virtual Storage/Enterprise Systems Architecture. A z/OS operating system environment that supports ESA/390.

MVS/ESA SP. An IBM licensed program used to control the z/OS operating system. MVS/ESA SP together with DFSMS compose the base MVS/ESA operating environment. See also *z/OS*.

mvslogin. A client command to connect your workstation to MVS.

mvslogout. A client command to break the connection between your workstation and MVS.

N

network. (1) An arrangement of nodes and connection branches. (2) A configuration of data processing devices and software connected for information interchange.

NFS. See *z/OS Network File System*.

NLM. Network Lock Manager.

nobody. A value of 65534 or -2 as userid in attributes file.

NSM. Network Status Manager.

null credentials. Null credentials are usually associated with diskless workstations where it is impossible to obtain identifying information because there is no local (that is, local to the workstation) repository of information.

O

object. A directory or file system.

object storage hierarchy. A hierarchy consisting of objects stored in DB2 table spaces on DASD, on optical or tape volumes that reside in a library, and an optical or tape volumes that reside on a shelf. See also *storage hierarchy*.

optical volume. Storage space on an optical disk, identified by a volume label. See also *volume*.

optical volume. Storage space on an optical disk, identified by a volume label. See also *volume*.

Operating System/2 (OS/2). An operating system used in IBM PC AT, PS/2, and compatible computers.

OSA2. Open System Adapter 2 (OSA2).

P

partitioned data set (PDS). A data set on direct access storage that is divided into partitions, called members, each of which can contain a program, part of a program, or data.

partitioned data set extended (PDSE). A system-managed data set that contains an indexed directory and members that are similar to the directory and members of partitioned data sets. A PDSE can be used instead of a partitioned data set.

PDS. See *partitioned data set (PDS)*.

PDS directory. A set of records in a partitioned data set (PDS) used to relate member names to their locations on a DASD volume.

PDSE. See *partitioned data set extended (PDSE)*.

permission code. A three-digit octal code, or a nine-letter alphabetic code, indicating the access permissions. The access permissions are read, write, and execute.

permission field. One of the three-character fields within the permissions column of a directory listing indicating the read, write, and run permissions for the file or directory owner, group, and all others.

PFS. Physical file system.

portable operating system interface (POSIX). Portable operating system interface for computer environments. An IEEE operating system standard, closely related to the UNIX system (software writing).

port. (1) An access point for data entry or exit. (2) A receptacle on a device to which a cable for another device is attached.

Portmapper. A server that converts RPC program numbers into port numbers acceptable to the protocol. This server must be running to make RPC calls.

POSIX . See *portable operating system interface (POSIX)*.

primary storage. A DASD volume available to users for data allocation. The volumes in primary storage are called primary volumes. See also *storage hierarchy*. Contrast with *migration level 1, migration level 2*.

protocol. (1) A set of semantic and syntactic rules that determines the behavior of functional units in achieving communication. (2) A specification for the format and relative timing of information exchanged between communicating parties.

PS. Physical sequential.

PTF. Program temporary fix.

Q

QSAM. See *queued sequential access method (QSAM)*.

queued sequential access method (QSAM). An extended version of the basic sequential access method (BSAM). Input data blocks awaiting processing or output data blocks awaiting transfer to auxiliary storage are queued on the system to minimize delays in I/O operations.

R

RACF. See *Resource Access Control Facility (RACF)*.

RBA. Relative byte address.

relative byte address (RBA). The displacement of a data record or a control interval from the beginning of the data set to which it belongs; independent of the manner in which the data set is stored.

relative record data set (RRDS). A type of VSAM data set whose records have fixed or variable lengths, and are accessed by relative record number.

remote procedure call (RPC). RPC is an independent set of functions used for accessing remote nodes on a network. Using RPC network services, applications can be created in much the same way a programmer writes software for a single computer using local procedure calls. The RPC protocols extend the concept of local procedure calls across the network, which means that distributed applications can be developed for transparent execution across a network.

Resource Access Control Facility (RACF). An IBM licensed program that is included in z/OS Security Server and is also available as a separate program for the z/OS and VM environments. RACF provides access control by identifying and verifying the users to the system, authorizing access to protected resources, logging detected unauthorized attempts to enter the system, and logging detected accesses to protected resources.

Resource Measurement Facility (RMF). An IBM licensed program or optional element of z/OS, that measures selected areas of system activity and presents the data collected in the format of printed reports, system management facilities (SMF) records, or display reports. Use RMF to evaluate system performance and identify reasons for performance problems.

rmdir. An AIX, UNIX, OS/2 and DOS command used to remove (delete) a directory.

RMF. See *Resource Management Facility (RMF)*.

root. The user name for the system user with the most authority.

RPC. Remote procedure call.

RRDS. Relative record data set.

S

SAF. System Authorization Facility (SAF).

SDSF. System Display and Search Facility (SDSF).

sequential file. A type of MVS file that has its records stored and retrieved according to their physical order within the file. It must be on a direct access volume.

server. (1) A functional unit that provides shared services to workstations over a network; for example, a file server, a print server, a mail server. (2) On a network, the computer that contains the data or provides the facilities to be accessed by other computers in the network. (3) A program that handles protocol, queuing, routing, and other tasks necessary for data transfer between devices in a computer system.

sharing. A term used in a computing environment to refer to utilizing a file on a remote system. It is done by mounting the remote file system, then reading or writing files in that remote system.

shell prompt. The character string on an AIX or UNIX command line indicating the system can accept a command (typically the $ character).

showattr. A client command used to display the values of the site, processing, and data set creation attributes.

SMF. See *system management facilities (SMF)*.

SMP/E. See *System Modification Program/Extended (SMP/E)*.

SMS. See *Storage Management Subsystem (SMS)*.

SPF. System productivity facility.

SRM. System Resources Manager (SRM).

SSF. Software support facility.

stale file handle. A file handle is stale when the file handle for a file or prefix is no longer valid.

stateless. A stateless server can function correctly without maintaining any protocol state information about any of its clients. The NFS protocol is intended to be as stateless as possible. This avoids complex crash recovery; a client just resends requests until a response is received. The stateless protocol is chosen to minimize the probability of data losses due to a server crash.

storage hierarchy. An arrangement of storage devices with different speeds and capacities. The levels of the storage hierarchy include main storage (memory, DASD cache), primary storage (DASD containing uncompressed data), migration level 1 (DASD containing data in a space-saving format), and migration level 2 (tape cartridges containing data in a space-saving format). See also *primary storage, migration level 1, migration level 2, object storage hierarchy*.

Storage Management Subsystem (SMS). A DFSMS facility used to automate and centralize the management of storage. Using SMS, a storage administrator describes data allocation characteristics, performance and availability goals, backup and retention requirements, and storage requirements to the system through data class, storage class, management class, storage group, and ACS routine definitions.

system management facilities (SMF). A component of z/OS that collects input/output (I/O) statistics, provided at the data set and storage class levels, which helps you monitor the performance of the direct access storage subsystem.

superuser. The user who can operate without the restrictions designed to prevent data loss or damage to the system (User ID 0).

System Modification Program/Extended. Basic tool for installing software changes in programming systems. It controls these changes at the element (module or macro) level, which helps protect system integrity.

T

tape volume. Storage space on a tape, identified by a volume label, which contains data sets or objects and available free space. A tape volume is the recording space on a single tape cartridge or reel. See also *volume*.

task control block (TCB). An MVS control block that describes an asynchronous process that is called a task.

TCB. See *task control block (TCB)*.

TCP/IP. See *Transmission Control Protocol/Internet Protocol (TCP/IP)*.

TIOT. Task input/output table.

transmission control block (TCB). An internal control block within the TCP/IP address space.

Transmission Control Protocol (TCP). A stream communication protocol that includes error recovery and flow control.

Transmission Control Protocol/Internet Protocol (TCP/IP). The two fundamental protocols of the Internet protocol suite. The abbreviation TCP/IP is frequently used to refer to this protocol suite. TCP/IP provides for the reliable transfer of data, while IP transmits the data through the network in the form of datagrams. Users can send mail, transfer files across the network, or execute commands on other systems.

TSO. Time-sharing option.

TSO/E. Time-sharing option extended.

U

UCB. Unit control block.

UDP. User datagram protocol.

UID. The DFSMShsm authorized-user identification in 1 to 7 characters. See also *user number (UID)*.

Unicode Standard. A universal character encoding standard that supports the interchange, processing, and display of text that is written in any of the languages of the modern world. It can also support many classical and historical texts and is continually being expanded. The Unicode Standard is compatible with ISO/IEC 10646.

UNIX. A highly portable operating system originally developed by Bell Laboratories that features multiprogramming in a multi-user environment. UNIX is implemented in the C language. UNIX was originally developed for use on minicomputers but has been adapted on mainframes and microcomputers. It is especially suitable for multiprocessor, graphics, and vector-processing systems.

UNIX authentication. Requires a client process to send credentials (the client's *machine-name*, *uid*, *gid*, and *group-access-list*) to the server.

user catalog. An optional catalog used in the same way as the master catalog and pointed to by the master catalog. It lessens the contention for the master catalog and facilitates volume portability.

User Datagram Protocol (UDP). A datagram protocol, without connections, that requires minimal overhead, but which does not guarantee delivery.

user number (UID). In the AIX operating system, a number that uniquely identifies a user to the system. It is the internal number associated with a user ID.

V

virtual storage access method (VSAM). An access method for direct or sequential processing of fixed and variable–length records on direct access storage devices. The records in a VSAM file can be organized in logical sequence by a key field (key sequence), in the physical sequence in which they are written on the file (entry sequence), or by relative record number.

volume. The storage space on DASD, tape, or optical devices, which is identified by a volume label. See also *DASD volume, optical volume, tape volume*.

VRA. Variable Recording Area (VRA).

VSAM. See *virtual storage access method (VSAM)*.

VSCR. Virtual storage constraint relief.

X

XDR. See *External Data Representation (XDR)*.

Z

z/OS. z/OS is a network computing-ready, integrated operating system consisting of more than 50 base elements and integrated optional features delivered as a configured, tested system. See also *MVS/ESA SP*.

z/OS Network File System. A base element of z/OS, that allows remote access to z/OS host processor data from workstations, personal computers, or any other system on a TCP/IP network that is using client software for the Network File System protocol.

z/OS UNIX System Services file system. In the z/OS UNIX environment, the collection of files and file management structures on a physical or logical mass storage device, such as a diskette, minidisk, or a disk drive. See also *hierarchical file system (HFS) data set*.

z/OS UNIX System Services (z/OS UNIX). The set of functions provided by the SHELL and UTILITIES, kernel, debugger, file system, C/C++ Run-Time Library, Language Environment, and other elements of the z/OS operating system that allow users to write and run application programs that conform to UNIX standards.

Index

Special characters

/etc/rpc file, client setup 12
/etc/services and tcpip.profile files, server setup 12
(<PROCNAME>) 130

Numerics

0037, EBCDIC code page 140
0C4 ABEND 227
0C4 protection exception 158
100003 3 udp 2049 nfs 220
2|3 parameter, vers client attribute 83
80A ABEND 109
878 ABEND 109

A

ABEND
  80A 109
  878 109, 252
  878 (80A) 253, 254
  X'0F3' 237
  X'806' 237
  X'A03' 237
  X'x13' 237
  X'x22' 237
  X'x37' 237
  X'x3E' 237
ABEND 878 252
ABEND 878 (80A) 253, 254
ABEND B37 268
ABEND D37 268
ABEND E37 268
ABEND0C4 227
ABENDxxx keyword 227
abytes value 321
access method services error message GFSA847I 265
access modes 184
access to remote files 67
accessibility 395
accessing MVS data sets
  command syntax for AIX 47
  command syntax for z/OS DFSMS 59
accessing MVS files 53
  changing attributes 52
  getting authorization 53, 70
  line terminators 28
  mount command 52, 65
  mvslogin command 64
  showattr command 76
accessing MVS files, displaying attributes 55
accessing z/OS files, showattr command 54
accessing z/OS UNIX System Services 38
acdirmax client attribute 79
acdirmax keyword 287
acdirmin client attribute 79
acdirmin keyword 287
acregmax client attribute 79
acregmax keyword 287
acregmin client attribute 79
acregmin keyword 287
active clients 271
active data sets, displaying 166
adjustment, packet size 81
afiles value 321
AIX client, Network File System 47
AIX command syntax 47
AIX commands 49
AIX, downloading client commands 150
AIX–retrieve commands 150
alias 295
alias name, message GFSA817I 263
alias names, MVS files 16
allocated buffers 104
allocated buffers, logical I/O processing 108
allocated storage, sidefile 115
allocating
  attributes data set 133
  checklist data set 138
  exports data set 134
  lock data sets 139
  mount handle data sets 139
altsym keyword 134, 135
altsym parameter, start command 159
American National Standard Code for Information Interchange (ASCII) 27, 39
ARC0075E DFSMShsm message 243
ASCII to EBCDIC conversion 93
Assembler header macro, reference 372
assignments, port 12
assortedparms 202
async attribtue 36
async attribute 37, 214
async processing attribute 98
asynchronous block input/output (I/O) daemons (BIOD) 80
asynchronous daemon task 291
atime 317, 320
attention messages, collecting 170
attrcaching attribute 219
attrcaching client attribute 80
attribute caching, consistency check 80
attribute statement syntax 134
attributes
  async 37, 98, 214
  attrtimeout 108, 249
  binary 93, 211
  blankstrip 28
  bufhigh 104, 247
  changing 52
  cln_ccsid(n), srv_ccsid(n) 100
  data set creation 52
  dsntype 89
  dsntype(pds) 18
attributes *(continued)*
  dsorg 89, 210
  dsorg(da) 17
  dsorg(ps) 16
  exports 7
  extlink 36
  fileextmap 30
  for z/OS UNIX System Services 87
  logicalcache 104
  logout 51
  maplower 23, 27
  maxtimeout 108, 250
  mintimeout 108, 109, 249
  mount point 55
  nfstasks 111, 213, 217
  noattrtimeout 93
  nofileextmap 30
  nomaplower 23, 27
  nomaxtimeout 108, 250
  noreadtimeout 108
  noretrieve 28
  nowritetimeout 108
  processing 52, 93
  readtimeout 109, 250
  recfm 90
  restimeout 53
  retieve 97
  retrieve 28
  retrieve(wait) 97, 213
  saf 51
  safexp 7, 51
  security 7, 51, 59, 175
  security(exports) 124, 126
  security(none) 126
  security(saf) 124, 361
  security(safexp) 124, 125, 361
  sidefile(dsname) 30
  site 104
  sync 98
  sync, async 36, 214
  text 93, 211
  trusted 119
  vol 249
  writetimeout 109, 250
  xlat 140, 141
  z/OS UNIX 36
attributes data set 133
attributes, data set 7
attrtimeout 213
attrtimeout attribute 249
attrtimeout message 250
attrtimeout processing attribute 93
attrtimeout site attribute 108, 109
attrtimeout timeout attribute 99
authentication error 64, 265
  restart processing 52
authentication, protocols 126
authority, trusted 119
authorization checking for z/OS UNIX files 37

© Copyright IBM Corp. 1991, 2006

Authorized Program Analysis Report (APAR) 225
authorizing file operations 125
authsys (system authentication) security 123
automatic class selection (ACS) 211
Automatic Class Selection (ACS) 92
automatic logout, recovering 53, 54, 70
automount facility 65, 133
aux parameter, space data set attribute 91
avg parameter, recordsize data set attribute 90

B

B37 ABEND 268
badcall statistic 198
badxid statistic 198
basic sequential access method (BSAM) 4, 67, 81
binary files 28
binary parameter, delim client attribute 81
binary processing attribute 93
binary processing mode 27
biod (block I/O daemons) 220
biod client attribute 80
blanks, handling of 28
blanks, trailing 93
blankstrip 295
blankstrip attribute 28
blankstrip processing attribute 93
blks data set attribute 88
BLKSIZE (block size) 209
blksize data set attribute 88
block I/O daemon (BIOD) 201, 215, 217
block I/O daemons (BIOD) 219, 220
block input/output (I/O) daemons (BIOD) 80
block size (BLKSIZE) 209
blocked I/O daemons (BIOD) 104
BPXVCLNY load module 158
BSAM (basic sequential access method) 67
BSAM, QSAM, and VSAM ESDS access to remote files 67
buffer client block writes, out of order 104
buffer size, write 83
buffer usage 214
buffers reclaimed, percentage 112
buffers, data 247
buffers, maximum size 104
buffers, remaining 108
bufhigh 214, 215, 217
bufhigh attribute 108, 219, 220, 247
bufhigh client attribute 80
bufhigh message 242, 247, 289
bufhigh site attribute 55, 104, 108
Bufhigh site attribute 77
bufhigh(n) limit 112

C

C function, fopen 288
C function, text1 289
C header macro, reference 365
cached buffers, disk 98
cached data, reading 112
cached file attributes, maximum/minimum 79
cachewindow 215, 217
cachewindow site attribute 55, 77, 104
caching 214
cachwindow() 27
cancel command 160
cancel mvsnfsc command 158
carriage-return newline (CRNL) 62
case_insensitive value 321, 322
case_preserving value 321, 322
cat command 18, 26, 28
cataloged data sets 5, 31
cccc parameter, recfm data set attribute 90
CCSID 297 66, 127
CCSID 437 66, 127
changing
  attributes 52, 65, 133
  mount handle data sets 169
  your MVS password 63
changing attributes for mount points 53
changing your z/OS password 50
channel status word, message GFSA876I 268
character data representation architecture (CDRA) 290
Character Data Representation Architecture (CDRA) 66, 126
checking UNIX permission bits 115
checklist attribute 125
checklist data set 8
CHECKLIST data set, sample 329
checklist site attribute 104
chmod command 24
chown_restricted value 321, 322
class_name parameter, storclas data set attribute 91
client attributes 79, 84
client commands, installing 147
client hanging 114
client log daemon 290
client log data set, setting up 379
client messages 281
client nfs 197
client program requirements 318
client rpc 197
client transport handle 278
client-server relationship 3
client, inactive time limit 108
client, noattrtimeout, noreadtimeout, nowritetimeout 108
client, z/OS NFS
  starting 157
  stopping 158
client, z/OS NFS introduction 3
client/server, user specified port range support 12
clients, tested 9
cln_ccsid client attribute 80
cln_ccsid message 255, 286
cln_ccsid parameter 11, 66, 67, 83, 127
cln_ccsid processing attribute 100, 101
cln_ccsid setting 67
cln_ccsid(n) processing attribute 94
clnt_call message 278
clnt_pcreateerror message 277
clnt_perror message 278
clntudp_create message 277
closing a data set 167
coded character set identifier (CCSID) 11, 80, 83, 98, 100, 255, 286
Coded Character Set Identifier (CCSID) 66, 86, 126
codes, reason 309
codes, return 299, 300, 302
collecting
  diagnostic messages 170
  usage statistics 130
command
  cancel 160
  cancel mvsnfsc 158
  cat 18, 26, 28
  chmod 24
  copy 19, 205
  cp 15, 205
  crnl2nl 62
  date 205
  dir 297
  force 158
  installing 63
  ln 40
  ls 18
  ls (UNIX) 94, 95, 314
  ls (UNIX), rm (UNIX) 94
  make (AIX or UNIX) 151
  make example 360
  mkdir 15, 17, 18, 33, 210
  modify 52, 53, 135, 164, 252, 260, 261, 270
  more 40
  mount 5, 15, 16, 28, 31, 48, 52, 53, 55, 59, 65, 125, 197, 209, 211, 213, 283, 284
    example 52
  mount (Windows or UNIX/Linux) 32
  mvslogin 47, 50, 51, 59, 64, 108, 124, 177, 276, 278, 298
  mvslogin (client) 147
  mvslogout 49, 50, 58, 63, 77, 177, 276
  mvslogout (client) 147
  net use 15
  netstat -s 198, 200
  nfsstat 61, 72, 197, 201
  nl2crnl 62
  rm 18, 41
  rmdir 18, 266
  showattr 16, 17, 37, 48, 50, 51, 54, 55, 60, 64, 76, 276, 277
  showattr (client) 147
  showmount 62, 75
  start 135, 159
  stop 159
  stop mvsnfsc 158
  time 205
  timex 205
  touch (AIX or UNIX) 151

command *(continued)*
TSO ALLOCATE 67
TSO HELP MOUNT 59
TSO HELP UNMOUNT 59
TSO MOUNT 66, 67, 127
umount 49, 57
umount, client 134
unmount 59, 71, 271, 284
vi 16, 26
command syntax
AIX user 47
comment symbol 159
comments from the requestor 363
commit parameter 37
commit procedure 27
compatibility, POSIX 36
component identification keyword 226
component trace
command for starting 161
concurrent write, PDSE 19
configuring the Network File System 119, 133
configuring the server
attributes data set 134
exports data set 134
mount handle data sets 134
console message (<PROCNAME>) 130
contiguous port range 12
control files 7
control timeout, HFS vnode token 106
controlling access to data sets 134
convserv client attribute 80
convserv processing attribute 103
values 128
customized conversion 128
enforced subset conversion 128
Language Environment-Behavior conversion 128
Modified Language Environment-Behavior conversion 128
roundtrip conversion 128
CONVXLAT utility 141
cookie verifier checking 96
copy command 19
cp command 15
cr parameter, delim client attribute 81
creating conventional MVS data sets 6
creating external link 40
creating VSAM files 19
credentials allocation failure, message GFSA827I 264
credentials, UNIX-style 126
crnl2nl command 62
ctime 317
ctrace
operand of the START mvsnfs command 159
customization, planning for 119
customized conversion 128
Customized conversion 103
customizing exit routines 173
customizing the translation table 140
cyls data set attribute 88

D

D37 ABEND 268
DASD volume, store data set 92
data access/creation commands 26
data buffers, memory blocks 247
data class, specifing 88
data conversion 83
data conversion (delim, filedata) 67
data conversion, crnl2nl, nl2crnl 66, 126
data conversion, delim, filedata 67
data management error message 268
data set
attributes 6, 7
binary and text formats 93
cataloged 5
cataloged, organizations supported 31
creation attributes 88
definition 5, 30
direct access 95
exports 7
MVS 186
non-system-managed 95
organization 89
PDS 95
readdirtimeout site attribute 114
record format and characteristics 90
record size (VSAM) 90
release unused space 91
released after a read 96
released after a write 98
rules for file extension mapping 97
serialization 34
share options (VSAM) 91
spanned data set attribute (VSAM) 91
store 92
structure 88
unit create 92
data set organization (DSORG) 210
data set type (DSNTYPE) 210
data set, side file 274
data transmitted, write request 98
data types, mixed set 96
data_sync parameter 37
databufferpoolsize 202
datacaching attribute 84, 219, 287
datacaching client attribute 81
DATACLAS 211
dataclas data set attribute 88
DCB (data control block) parameters 88
DDNAME EXPORTS 267
DDNAME NFSATTR 248
deallocating a data set 167
debugging, collecting messages for 170
default attributes, displaying 55
default attributes, overriding 52, 65
defaults, displaying 54, 76
delay parameter 71
delay write, disk blocks 81
delaywrite attribute 219, 220
delaywrite client attribute 81
delaywrite keyword 287
deleting
entries from mount handle data sets 169
migrated files 97
deleting *(continued)*
tasks 252
delim
client attribute 81
parameter 67, 288
denyrw site attribute 105
development toolkits 153, 155
DFSMS command syntax 59
DFSMShsm 97
DFSMShsm, recall or delete migrated files 97
dhcp site attribute 105
diagnosing problems 170
diagnosis and reporting 225
diagnosis, using keywords 225
diagnostic aids 236
dir command 297
dir data set attribute 89
direct access (DA) 15, 314, 317
creating 17
data set 89
directory
client commands 89
record number 89
disability 395
disable lookup lookaside caching attribute 81
discarded, data set 114
disconnecting a mount point 57, 71
disk blocks
maximum 81
read ahead 82
disk, physical block 88
dispatching priority 202
displaying
client statistical information 72
default and mount point attributes 76
default attributes 77
mount points and active data sets 166
remote server mount information 75
site and mount point attributes 16
status of active subtasks 168
displaying default and mount point attributes 54, 55
displaying default attributes 55
DOC SCnnnnnnnn keyword 228
domain name server 263
Domain Name Server 159
double byte character set (DBCS) 66, 127, 140
download, commands to the client 147
downloading
source code to the client 359
dsname parameter, sidefile processing attribute 97
DSNTYPE (data set type) 210
dsntype data set attribute 89
dsntype message 251
dsntype(library) attribute 18
dsntype(pds) attribute 18
DSORG (data set organization) 210
dsorg data set attribute 89
dsorg message 249
dsorg(da) attribute 17
dsorg(da) parameter 17

dsorg(ps) example, attribute 16
dsps
  operand of the START mvsnfs command 159
dtpref value 321, 322
Dump Analysis and Elimination (DAE) 237
dump data set, abnormal end 235
duration parameter 71
dynamic allocation failure 265, 266
dynamic IP addresses for clients 105
dynamic IP addressing 12
  server configuration 147
dynamicsizeadj attribute 219
dynamicsizeadj client attribute 81

E

E37 ABEND 268
EBCDIC code page 0037 140
EBCDIC to ASCII conversion 93
Electronic Technical Response (ETR) 235
encoded password 363
encoded user name 363
end-of-file mark (EOF) 62
end-of-line (EOL) terminator 270
end-of-line specifiers 28, 94
ending your z/OS session 58, 77
enforced subset conversion 128
Enforced Subset conversion 103
Enterprise Storage Server (ESS) 210
entry-sequenced data set (ESDS) 15, 19, 67
envar parameter 129
error messages 241
error messages, collecting 170
error, side file data set 274
ESDS (entry-sequenced data set) 67
ESS 2105 DASD 210
ESTAE exit routine 254
estae macro 291
etc/passwd 276, 278
executebitoff processing attribute 94
executebiton processing attribute 94
exit routines 173
export entry, return first character 107
export spanning pathnames 10
Export Spanning Pathnames 252
EXPORTFS command 275
exportfs operand 52, 166, 272
exporting a file system 125
exports data set 7, 298
exports list and SAF checking 124
extended binary-coded decimal interchange code (EBCDIC) 27, 39
extended format data sets, exploiting 21
External Data Representation (XDR) 4
external link 40
externalized return code mapping, NFS version 2 protocol 302
externalized return codes, NFS version 2 protocol 299
externalized return codes, NFS version 3 protocol 299
externalized return codes, NFS version 4 protocol 300
extlink attribute 36
extlink processing attribute 94

F

fastfilesize 202, 211
fastfilesize processing attribute 95, 314
FAT file system 153
fbytes value 321
ffiles value 321
file attributes, NFS version 4 protocol 321
file creation attributes 88
file creation, data class 88
file creation, management class 90
file creation, storage class 91
file extension mapping 30
file naming conventions 16
file parameter, smf site attribute 116
file processing, override default translation table 99
file security 279
file security exit 181, 184
  request code 186
  return code 186
file size determination 24, 313
file size, number physical block buffers 108
file system 5, 30
file system attributes for MVS 34
file system size 42
file tag, new files 98
file tagging, Unicode Services active 101
file tagging, Unicode Services not active 103
file timeout 260
File Transfer Protocol (FTP) 147
file usage block, message GFSA881I 269
file_sync parameter 37
filedata parameter 67
fileextmap attribute 30
fileextmap processing attribute 95
fileidsize site attribute 105
files (MVS), saving 28
files, locking 43
files, source for client enabling commands 153
filesystype parameter 289
filesystype parmlib statement 132
fill internal buffers, read processing 114
fixed-blocked file format 88
fixed-length text file, record 93
flushctr operand 170
flushing partial records, message GFSA811I 262
flushlog operand 170
fn_delimiter site attribute 105
force command 158
freeze operand 166
freeze=off operand 53, 270
freeze=offhfs operand 53
freeze=on operand 53, 270
freeze=onhfs operand 53
French EBCDIC format (CCSID 297) 66, 127
FSF_CANSETTIME, properties 321
FSF_HOMOGENEOUS, properties 321
FSF_LINK, properties 321
FSF_SYMLINK, properties 321
fsinfo, static file system 321, 322
fsstat, dynamic file system 321

G

Get Attribute 211, 213
getattr 250
getattr call 198
getattr operation 93
GFSA990I 184
GFSA991E 184
GFSASSMF macro 365
GFSAUDSA user storage block 173
GFSAUSMF macro 372
GID to group name, mapping 362
global exit block (GXB) 181
Global Exit Block (GXB) 174, 178
grace period for reclaiming locks 44
granting access to data sets 134
group name to GID, mapping 362
group number (GID) 24, 38

H

hanging, client 114
HDZ11VC, HDZ11VS 226
HDZ11VS 158
hfs parameter, security site attribute 115
hfs site attribute 106
HFS site attribute 37
hfs_path parameter, public site attribute 113
HFS() keyword 113
hfsfbtimeout site attribute 106
hfssec site attribute 107
hierarchical file system (HFS) 4, 35, 40, 41, 100, 211, 214, 217, 251, 260, 261
high-level qualifier (HLQ) 5
high-level qualifiers 30
hostname
  mvslogin command 64

I

IBMLink/Service, diagnosis using 225
IDCAMS error 265
inactive time limit 108
incorrout keyword 228
increase network bandwidth 200
indexed, data set 89
informational messages, collecting 170
input errors 225
installation default settings, overriding 52
installation exits 126
installation parameters 85
installing commands 50, 63
installing the mvslogin, mvslogout, and showattr commands 147, 359
Internet Protocol (IP) address 7
  dynamic 12, 105, 147
internet protocol version 6
  SMF records for 131
interprocess communication (IPC) queue 257

invarsec value 321
ISHELL utility 36
ISO 8859-(ASCII) 140

J

job control language (JCL) 88, 267

K

Kerberos
  authentication with NFS v4 120
  key distribution center 391
key-sequenced data set (KSDS) 15, 19, 244
keyboard 395
keys data set attribute 89
keys, length and offset 89
keyword
  altsym 135
  security 136
keyword, using 225
krb5 parameters
  in the exports data set 136
  on hfssec site attribute 107
  on mvssec site attribute 110
  on pubsec site attribute 113
KSDS (key-sequenced data set) 67

L

Language Environment-Behavior conversion 103, 128
large format data sets, accessing 21
leadswitch site attribute 107
leasetime for locks 44
leasetime site attribute 107
legacy_path parameter, public site attribute 113
len parameter, keys data set attribute 89
lf end-of-line specifier 28
lf parameter, delim client attribute 81
lfcr parameter, delim client attribute 81
limit, storage 80
limitations, PDS 19
limiting access to data sets 134
line terminators 28
link request 34
linkmax value 321, 322
list operand 166
list=dsnames operand 271, 272
list=mounts operand 261, 271, 272
listing lock holders for a file 44, 170
listlock operand 44, 170
ln command 40
local host name 363
local mounted file system 80, 98
localpath 48
lock data sets 139
lock holders
  listing 44, 170
lock time limit 44
lock time, control 107
lockd protocol 327
locking and access control 43
log data set 8
log data set, setting up 379
log data sets
  NFSLOG1 164, 170
  NFSLOG2 164, 170
log data sets, NFS client 380
log data sets, NFS server 379
log=msglevel operand 171
logical file system (LFS) 132
logical I/O processing, cache windows 108
logical I/O, buffer client block writes 104
logicalcache 215, 217
logicalcache attribute 247
logicalcache site attribute 55, 77, 104, 108
login and logout with PCNFSD 361
login exit routine 175
logon and file security 279
logout
  logout, forced 175
logout attribute 51
logout attribute, maxtimeout site attribute 108
logout site attribute 55, 77, 108
logout, recovering from automatic 53, 54, 70
LookAt message retrieval tool xv
lookup 93, 198
Lookup 213
lookup request 9, 113, 252
loop keyword 228
low system usage 200
lower case to upper case 95
lrecl data set attribute 90
ls (UNIX) command 94, 95, 314
ls command 18

M

m parameter, nfstasks site attribute 111
m parameter, restimeout site attribute 114
macro
  STOW 262, 263
make (AIX or UNIX) command 151
make command example 360
management class, file creation 90
mapfile 273, 274
mapfile operand 167
mapleaddot processing attribute 95
maplower attribute 23, 27
maplower processing attribute 95
mapped keyword processing attribute 100
mapped processing attribute 96
mapped, side files 273
mapping, file extension 97
mapping, lower case file names 95
mapping, procedure 12 362
mapping, return codes 302
mapping, UID, user name, GID, group name 362
max parameter, recordsize data set attribute 90
maxfileproc parameter 292
maxfilesize value 321
maximum disk blocks 81
maximum size, allocated buffers 108
maximum transmission unit (MTU) 201
maxrdforszleft 215
maxrdforszleft site attribute 55, 77, 108
maxsockets 257
maxtimeout 213
maxtimeout attribute 250
maxtimeout message 250
maxtimeout site attribute 55, 77, 108, 109
member_name parameter, xlat processing attribute 99
memory blocks, data buffers 247
message log data set, setting up 379
message retrieval tool, LookAt xv
message, console 130
messages, client
  GFSC100E 282
  GFSC101E 282
  GFSC102E 283
  GFSC103E 283
  GFSC105E 283
  GFSC106E 283
  GFSC107E 283
  GFSC110E 283
  GFSC200E 283
  GFSC201E 284
  GFSC202E 284
  GFSC203E 284
  GFSC204E 284
  GFSC205E 284
  GFSC206E 284
  GFSC207E 284
  GFSC208E 284
  GFSC209E 285
  GFSC210E 285
  GFSC211E 285
  GFSC212E 285
  GFSC213E 285
  GFSC214E 285
  GFSC300E 285
  GFSC301E 286
  GFSC302E 286
  GFSC303E 286
  GFSC304E 286
  GFSC305E 286
  GFSC307E 286
  GFSC308E 286
  GFSC309E 287
  GFSC310I 287
  GFSC311I 287
  GFSC312I 287
  GFSC313I 287
  GFSC315E 287
  GFSC317E 287
  GFSC318E 287
  GFSC319E 288
  GFSC320E 288
  GFSC500I 288
  GFSC501I 288
  GFSC502E 288
  GFSC503E 288
  GFSC504I 288
  GFSC505E 288
  GFSC506E 289
  GFSC700I 289

messages, client *(continued)*
GFSC701I 289
GFSC702I 289
GFSC703E 289
GFSC704E 289
GFSC705E 289
GFSC707E 289
GFSC708E 290
GFSC709E 290
GFSC710E 290
GFSC711E 290
GFSC712E 290
GFSC713E 290
GFSC714E 290
GFSC715E 290
GFSC716I 291
GFSC721E 291
GFSC722E 291
GFSC723E 291
GFSC724E 291
GFSC725E 291
GFSC726E 292
GFSC727W 292
GFSC728E 292
GFSC729W 292
GFSC840I 292
GFSC841E 292
GFSC843E 292
GFSC845I 293
GFSC846E 293
GFSC847E 293
GFSC848E 293
GFSC849E 293
GFSC850E 293
GFSC854I 293
GFSC855E 294
GFSC856E 294
GFSC858E 294
messages, client platform (AIX)
Cross device link message 295
Directory Not Empty message 295
File exists message 295
File Name Too Long message 295
I/O Error (with possible system programmer response) 295
I/O Error (with possible user response) 295
Invalid 295
Is a directory 296
No space left on device 297
No such device 296
No such file or directory 297
Not a directory 297
Not Owner 297
Permission denied 297
Read Only File System 298
server name not responding still trying 296
Stale NFS File Handle 298
value too large 298
Weak Authorization 298
messages, display in mixed/upper case 109
messages, server
GFSA320I (procname) 242
GFSA322I (procname) 242
GFSA325I 242
messages, server *(continued)*
GFSA326E 243
GFSA328I (procname) 243
GFSA329I 160
GFSA329I (procname) 243
GFSA330I 160
GFSA330I (procname) 243
GFSA331E (procname) 243
GFSA333I (procname) 243
GFSA334E (procname) 243
GFSA335E (procname) 244
GFSA335I 244
GFSA336E (procname) 244
GFSA337E 244
GFSA338E (procname) 245
GFSA339I 245
GFSA340I 245
GFSA344I (procname) 245
GFSA345E (procname) 245
GFSA346I 246
GFSA347I 246
GFSA348I 159
GFSA348I (procname) 246
GFSA349I 246
GFSA350E 246
GFSA351E 246
GFSA352E 247
GFSA360I 247
GFSA361I 247
GFSA362I (procname) 247
GFSA363I (procname) 247
GFSA364I (procname) 247
GFSA400I 248
GFSA401I (procname) 248
GFSA402I (procname) 248
GFSA403I 249
GFSA403I (procname) 248
GFSA404I 248
GFSA405I 248
GFSA406I 248
GFSA407I 248
GFSA408I 249
GFSA409I 249
GFSA411I 249
GFSA412I 249
GFSA413I 249
GFSA414I 249
GFSA415I 249
GFSA416I 249
GFSA417I 249
GFSA418I 250
GFSA419I 250
GFSA420I 250
GFSA421I 250
GFSA422I 250
GFSA423I 250
GFSA424I 250
GFSA425I 251
GFSA426I 251
GFSA429I 251
GFSA430I 251
GFSA431I 100, 251
GFSA434I (procname) 251
GFSA435E 251
GFSA436I 251
GFSA437I 251
GFSA438I 252
messages, server *(continued)*
GFSA439I 252
GFSA440I 252
GFSA441E 252
GFSA450I 252
GFSA451I 252
GFSA452I 252
GFSA453E (procname) 252
GFSA454I (procname) 253
GFSA455E 253
GFSA456E 253
GFSA459E 253
GFSA460E 253
GFSA461E (procname) 253
GFSA470I (<PROCNAME>) 237
GFSA470I (procname) 254
GFSA471I (<PROCNAME>) 237
GFSA471I (procname) 254
GFSA472I (procname) 254
GFSA473I (procname) 254
GFSA480I 254
GFSA481I 254
GFSA482I 255
GFSA483E 255
GFSA483I 255
GFSA484E 255
GFSA485I 255
GFSA486I 255
GFSA487E 255
GFSA501I 256
GFSA502I 256
GFSA550I 256
GFSA554I (procname) 256
GFSA555I 256
GFSA556I (procname) 256
GFSA557I 256
GFSA558I (procname) 257
GFSA559I (procname) 257
GFSA564I (procname) 257
GFSA566I (procname) 257
GFSA727E 257
GFSA728E 257
GFSA729E (procname) 258
GFSA730I (procname) 258
GFSA731I 258
GFSA733I 258
GFSA734E 258
GFSA735I 258
GFSA736E 259
GFSA737I 259
GFSA738E 259
GFSA739I 259
GFSA750I (procname) 260
GFSA751I 168
GFSA751I (procname) 260
GFSA752I (procname) 260
GFSA753I 168
GFSA753I (procname) 260
GFSA754I (procname) 260
GFSA770I 260
GFSA771I 166
GFSA771I (procname) 260
GFSA772I 166
GFSA772I (procname) 260
GFSA776I 261
GFSA777I 261
GFSA781I 168

messages, server *(continued)*
GFSA782I (procname) 261
GFSA784I 261
GFSA786I 261
GFSA787I 261
GFSA791I 261
GFSA792I 262
GFSA793I 171, 262
GFSA794I 171, 262
GFSA801I 262
GFSA802E 262
GFSA811I 262
GFSA812I 262
GFSA814I 263
GFSA815I 263
GFSA816I 263
GFSA817I 263
GFSA818I (procname) 263
GFSA819I 264
GFSA820I 264
GFSA821I 264
GFSA822I 264
GFSA823I (procname) 264
GFSA827I 264
GFSA829I 265
GFSA832I 265
GFSA833I 265
GFSA840I 265
GFSA842I 265
GFSA843I 265
GFSA847I 265, 272
GFSA848I 266
GFSA849I 266
GFSA853I 265, 266
GFSA854I 265, 266
GFSA858I 266
GFSA859I 266
GFSA860I 266
GFSA862I 267
GFSA863I 267
GFSA864I (procname) 267
GFSA865I (procname) 267
GFSA866I (procname) 267
GFSA867I (procname) 267
GFSA868I (procname) 267
GFSA869I (procname) 267
GFSA871I 268
GFSA876I 268
GFSA877I 268
GFSA878I (procname) 268
GFSA879I (procname) 269
GFSA881I 269
GFSA886I 269
GFSA895I 269
GFSA896I 269
GFSA897I 269
GFSA898I 270
GFSA899I 270
GFSA900I 168
GFSA900I (procname) 270
GFSA901I 166
GFSA901I (procname) 270
GFSA902I 166
GFSA902I (procname) 270
GFSA903I 168
GFSA903I (procname) 270
GFSA904I (procname) 270
messages, server *(continued)*
GFSA905I 270
GFSA908I (procname) 271
GFSA909E (procname) 271
GFSA910I 166
GFSA910I (procname) 271
GFSA911I (procname) 271
GFSA912I 167
GFSA912I (procname) 271
GFSA913I (procname) 271
GFSA914I 167
GFSA914I (procname) 271
GFSA915I 167
GFSA915I (procname) 271
GFSA916I 169
GFSA916I (procname) 271
GFSA917I 169
GFSA917I (procname) 271
GFSA918I (procname) 272
GFSA919I (procname) 272
GFSA920I (procname) 272
GFSA921I (procname) 272
GFSA922I (procname) 272
GFSA923I (procname) 272
GFSA924I (procname) 272
GFSA925I (procname) 272
GFSA926I (procname) 273
GFSA927I (procname) 273
GFSA928I 273
GFSA929I 273
GFSA930I (procname) 273
GFSA931I (procname) 273
GFSA932I (procname) 273
GFSA933I (procname) 273
GFSA934I (procname) 273
GFSA935I (procname) 274
GFSA936I (procname) 274
GFSA937E 274
GFSA938I 274
GFSA939E 274
GFSA940E 274
GFSA941I 275
GFSA942I 275
GFSA943E (procname) 275
GFSA944I 275
GFSA945I 275
GFSA946I 275
GFSA947I 275
GFSA948I 276
GFSA949I 276
GFSA950I 276
GFSA951I 276
GFSA952I 276
GFSA953I 276
GFSA954I 276
GFSA955I 276
GFSA956I 276
GFSA957I 276
GFSA958I 276
GFSA959I 276
GFSA960I 276
GFSA961I 277
GFSA962E 277
GFSA963E 277
GFSA964I 277
GFSA965I 277
GFSA966I 277
messages, server *(continued)*
GFSA967I 277
GFSA968I 277
GFSA969I 277
GFSA970I 277
GFSA971I 278
GFSA972I 278
GFSA973A 278
GFSA974A 276, 278
GFSA975A 278
GFSA975I 276
GFSA976I 278
GFSA977I 278
GFSA978I 278
GFSA979E 278
GFSA980E 278
GFSA981I 278
GFSA982E 278
GFSA983E 279
GFSA984E 279
GFSA985E 279
GFSA986I 279
GFSA987I 279
GFSA990I 177, 279
GFSA991E 177
GFSA991E (procname) 279
GFSA996E (procname) 279
GFSA997I 280
GFSA998I 280
mgmt_class_name parameter, mgmtclas data set attribute 90
mgmtclas data set attribute 90
migrated files 97
migrated files, recall or delete 97
mintasks site attribute 109
mintimeout 213
mintimeout attribute 249
mintimeout message 249, 250
mintimeout site attribute 55, 77, 108, 109
mixcase site attribute 109
mixed case, message display 109
mixed set of data types 96
mkdir (UNIX) command 17
mkdir command 15, 18, 33, 210
mknod request 34
Modified Language Environment-Behavior conversion 103, 128
modify command 52, 53, 131, 135, 164, 252, 260, 261, 270
modify command, operands of 159
modifying
attributes 133
checklist data set 138
mount handle data sets 169
modifying file attributes 52, 65
monitored lock 43
monitored locks 43
more command 40
mount and umount with PCNFSD 361
mount command 5, 15, 16, 28, 31, 48, 53, 59, 65, 125, 197, 209, 211, 213, 283, 284
changing site and mount point attributes 52
overriding default attributes 52
mount emulation if NFS V4 96
mount handle data set 8, 139
Mount Handle Database failed 262

mount operation, retry 82
mount point 88, 96, 104, 114, 275
  command 33
  creating a mount point for a PDS 18
  definition 31
  disconnecting 57, 71
  multiple 33
mount point, access 18
mount point, changing attributes 53
mount point, definition 5
mount points, saving 53, 70
mount processing parameters 84
mount requests 106
mount tag option 86
mount, umount, login, logout 361
mounted file space 271
mounted, local 80
MSGGFShnnnt keyword 227
mtime 317, 320
MTU (Packet Size) 238
multiple byte character set (MBCS) 140
multiple data set creation attributes, specifying 88
multiple NFS servers 160
MVS
  password and user ID 63
MVS file attributes 87
MVS files 54
  attributes 79
  changing attributes 52
  line terminators 28
  mount command 52, 65
  mvslogin command 64
  mvslogout command 58, 77
  saving 28
  selecting format 23
  showattr command 76
  umount command 57
  unmount command 71
MVS files, alias names 16
MVS files, creating 15
MVS files, creating PDSs and PDSEs 17
MVS files, direct access 17
MVS files, physical sequential 16
mvs parameter, security site attribute 115
MVS programs 33
MVS, accessing 49
MVS/ESA 36
mvslogin 119
mvslogin command 47, 50, 51, 59, 63, 108, 124, 147, 276, 278, 298, 361
  authentication errors 64
  syntax for clients 64
mvslogout 119
mvslogout command 49, 50, 58, 63, 147, 276
MVSMNT processing attribute 96
MVSNFS startup procedure 138
mvssec site attribute 110

N

n parameter, acdirmax client attribute 79
n parameter, acdirmin client attribute 79
n parameter, acregmax client attribute 79
n parameter, acregmin client attribute 79
n parameter, attrtimeout processing attribute 93
n parameter, biod client attribute 80
n parameter, bufhigh site attribute 104
n parameter, cachewindow site attribute 104
n parameter, cln_ccsid processing attribute 100
n parameter, delaywrite client attribute 81
n parameter, logicalcache site attribute 108
n parameter, logout site attribute 108
n parameter, maxrdforszleft site attribute 108
n parameter, maxtimeout site attribute 108
n parameter, mintimeout site attribute 109
n parameter, nfstasks site attribute 111
n parameter, percentsteal site attribute 112
n parameter, readahead client attribute 82
n parameter, readaheadmax site attribute 114
n parameter, resize client attribute 82
n parameter, restimeout site attribute 114
n parameter, retrans client attribute 82
n parameter, retry client attribute 82
n parameter, sfmax site attribute 115
n parameter, srv_ccsid processing attribute 98, 100
n parameter, timeo client attribute 83
n parameter, writetimeout processing attribute 98
n parameter, wsize client attribute 83
name_max value 321, 322
namesrv 159
naming conventions, files 16
Native ASCII environment support 100
native ASCII NFS client support 85
native ASCII support 11
native path 10
netstat 200
netstat command 198, 200
Network File System Protocol Specification, RFC 1094 282, 283
Network Lock Manager (NLM) 43
Network Status Monitor (NSM) 44
new files, file tag 98
newline (NL) 62
NFS client
  starting 157
  stopping 158
NFS clients, non-z/OS based 85
nfs link command 15
NFS protocol 9
  attributes for NFS server 9
NFS protocol attributes, z/OS NFS server support for 323
NFS servers, non-390 based NFS clients 101
NFS version 2 protocol, externalized return code mapping 302
NFS version 2 protocol, externalized return codes 299
NFS version 3 protocol, externalized return codes 299
NFS version 3 protocol, return codes 304
NFS version 4 file attributes 321
NFS version 4 protocol, externalized return codes 300
NFS version 4 protocol, return codes 306
nfsattr 297
NFSATTR, site attribute file 273
NFSC_001 72
NFSERR_STALE error response 261
NFSLOG1 log data set 164, 170
NFSLOG2 log data set 164, 170
nfsstat 221
nfsstat -m 294
nfsstat -m /mountpoint command 220
nfsstat command 61, 72
  nfsstat -c 197, 201
  nfsstat -in 201
  nfsstat -s 201
  nfsstat -z 197
NFSTARB files, downloading 147
nfstasks 213, 217
nfstasks attribute 217
nfstasks message 252
nfstasks site attribute 55, 77, 111
NFSXLAT DD statement 279
nl parameter, delim client attribute 81
nl2crnl command 62
nlm site attribute 112
no space condition 273
no_trunc value 321, 322
noattrtimeout message 250
noattrtimeout processing attribute 93
noattrtimeout site attribute 108
noblankstrip processing attribute 93
nochecklist site attribute 104
nodenyrw site attribute 105
nodhcp site attribute 105
nofastfilesize 202, 211, 213
nofastfilesize processing attribute 95, 315
nofileextmap attribute 30
nofileextmap processing attribute 95
noleadswitch site attribute 107
nomapleaddot processing attribute 95
nomaplower attribute 23, 27
nomaplower processing attribute 95
nomaxtimeout message 250
nomaxtimeout site attribute 108
non-monitored lock 43
non-monitored locks 43
none parameter, smf site attribute 116
nonindexed, data set 89
nonlm site attribute 112
nonspanned data set attribute 91
nopcnfsd site attribute 112
nopublic site attribute 113
noreadtimeout message 250
noreadtimeout processing attribute 96
noreadtimeout site attribute 108
norec878 site attribute 112
noretrieve 213

noretrieve attribute 28
noretrieve processing attribute 97
norlse data set attribute 91
not loaded, side file data set 274
notag processing attribute 98
noudpqueuelimit 202
nowait parameter, retrieve processing attribute 97
nowritetimeout message 250
nowritetimeout processing attribute 98
nowritetimeout site attribute 108
number of active clients 271
number of physical block buffers 108
number representation 28
numbered, data set 89

O

o parameter, nfstasks site attribute 111
obtaining an MVS password 50, 63
OEDIT editor, UNIX 19
OEMVS311 translation table 39, 141
off parameter, keys data set attribute 89
off parameter, smf site attribute 116
on parameter, smf site attribute 116
operands of the MODIFY command 159
operating the client 157
operating the client, stopping 158
operating the NFS NSM, NLM 161
operating the server 158
operating the server, stopping 160
ordering out-of-sequence data 215
out of order, buffer client block writes 104
override, default translation table 99
overriding default attributes 52, 65
overview of Network File System 3

P

packet size adjustment, remote procedure call (RPC) 81
parameter
  commit 37
  data_sync 37
  delay 71
  delim 67
  duration 71
  envar 129
  file_sync 37
  retrans 213
  stable 37
  unstable 37
  vers(x) 9
  xlat 67
parameter list contents 177
parameters
  region 242
  xlat() 255
parsing failure in attributes data set, message GFSA403I 248
partitioned data set (PDS) 6, 15, 18, 30, 33, 140, 141, 210, 314, 318, 320
partitioned data set extended (PDSE) 6, 17, 30, 33, 140, 141, 210, 318, 320
password (MVS) 63, 124
password (z/OS) 50
pathconf, retrieve POSIX information 321, 322
PCNFSD authentication 361
PCNFSD authentication request 51
pcnfsd site attribute 112
PCNFSD, authentication 119
PCNFSD, login, logout, mount, unmount 361
PCNFSD, version 1 protocol 362
PCNFSD, version 2 protocol 362
PDS directory, number of records 89
PDS or PDSE, removing 18
PDS or PDSE, timing out while writing 19
PDS or PDSE, update or extend 19
PDS or PDSE, wildcard copy restriction 19
PDS, creating 18
PDS, limitations 19
PDS, time stamps 318
PDSE or PDS, create 89
PDSE, concurrent write 19
PDSE, creating 17
PDSE, time stamps 318
percentage, buffers reclaimed 112
percentsteal 214, 215
percentsteal site attribute 55, 77, 104, 112
perfm keyword 229
perform mapping 362
permission bits, checking UNIX 115
permission denied message 51, 64
permissions, security 181
physical block buffers, number 108
physical block, maximum length 88
physical sequential 6
physical sequential (PS) 317
physical sequential (PS) files 15
physical sequential files, creating 16
planning for installation, security measures 119
port assignments 12
port range, user specified 143
port registration failure, message GFSA554I (procname) 256
portable operating system interface (POSIX) 3, 36
porting commands for clients 153
portmap 157, 159
portmapper 10, 239, 257
prefix parameter, hfs site attribute 106
prim parameter, space data set attribute 91
procedure 12, mapping 362
procedure 13, user authentication 363
processing attributes 93, 100
PROCLIB updates 132
protecting the file system
  with NFS v4 security 120
  with the security site attribute 123
protecting your programs and files 119
protecting your z/OS UNIX System Services files 38
proto attribute 10
proto client attribute 82
proto, mount request 285
protocol attributes, z/OS NFS server support for 323
protocol, PCNFSD version 1 362
protocol, PCNFSD version 2 362
PTFs, APARs, search 235
public file handle 113, 115
public keyword 9, 251, 264
public parameter, security site attribute 115
public site attribute 113
pubsec site attribute 113

Q

QSAM (queued sequential access method) 67
qualifiers 30
qualifiers, high-level 5
queued sequential access method (QSAM) 4, 67, 81

R

R1VS, R1VC 226
RACF (resource access control facility) 119, 124
RACROUTE, remote client 120
RAMAC3 210
random access to files 26
rdrverf processing attribute 96
read ahead, maximum disk blocks 82
read buffer size, set 82
read statistic 198
readahead attribute 219, 220
readahead buffers 112
readahead client attribute 82
readahead keyword 287
readaheadmax 215
readaheadmax site attribute 55, 77, 114
readdir on root 267
readdirplus request 34, 72
readdirtimeout site attribute 114
reading out of cached data 112
readlink request 34
readtimeout 213
readtimeout attribute 250
readtimeout message 249
readtimeout processing attribute 96
readtimeout site attribute 109
readtimeout timeout attribute 99
reason code
  rsnc 282
reason code reasoncd 290
reason codes 309
reasoncd reason code 290
rec878 site attribute 112
recalling migrated files 97
recfm data set attribute 90
recfm message 251
recno, write 269
record boundaries 28
record format (RECFM) 210
record length (LRECL) 209
record number, directory of a PDS 89
record type-42 (subtype 7, subtype 8) 131

record, fixed-length text file 93
records, locking 43
recordsize data set attribute 90
recs data set attribute 88
reduce data transfer 200
reference
  Assembler header macro 372
  C header macro 365
region parameter 242
region size 289
relative record 19
relative record data set (RRDS) 15, 19
release level keyword 226
release operand 167, 272
release=dsname(member) operand 271
releasing locks 44, 167
remote mount function 53
remote procedure call (RPC) 4, 73, 81, 82, 83, 158, 256, 278, 282, 292, 298
remove request 41
removing entries from mount handle data sets 169
rename failure, messages GFSA814I and GFSA815I 263
request code
  file security exit 186
  login exit 186
requests, rm, rmdir, mv 16
resize client attribute 82
resource access control facility (RACF) 24, 120
Resource Access Control Facility (RACF) 38, 261, 289, 297
Resource Measurement Facility (RMF) 206
restart processing 64
restarting the z/OS NFS client 157
restimeout attribute 53
restimeout site attribute 8, 114
restricting access to data sets 134
resuming mount processing 166
RETAIN, search 236
retc return code 283
retrans client attribute 82
retrans option 83
retrans parameter 213
retransmit, remote procedure call (RPC) 82
retrieve attribute 28
retrieve attributes 99
retrieve processing attribute 97
retrieve(nowait) 213
retrieve(wait) 213
retrieve(wait) attribute 213
retrieving client commands 150
retry client attribute 82
retry mount operation, specify 82
return code
  file security exit 186
  login exit 186
  returncd 282
  rsnc 290
return code 132 269
return code retc 283
return code returncd 283, 284, 290
return codes 299, 300, 302
return codes, NFS version 3 protocol 304
return codes, NFS version 4 protocol 306
returncd return code 282, 283, 284, 290
reuse option 19
RFC 2055 9
rlse data set attribute 91
rm (UNIX) command 94
rm command 18, 41
rmdir command 18, 266
root directory 4, 35
roundtrip conversion 128
Roundtrip conversion 103
RPC 1094 282, 283
rpcbind client attribute 82
rpcinfo -p servername command 220
RRDS (relative record data set) 67
RS/6000, downloading client commands 150
rsize 209, 221
rsnc reason code 282
rsnc return code 290
rtmax value 321, 322
rtmult value 321, 322
rtpref value 321, 322

S

SAF 8
saf attribute 51
saf checking, bypass 104
SAFEXP 8
safexp attribute 51
SAM striped files 20
sample CHECKLIST data set 329
samples
  exports data set 349
  startup procedures z/OS NFS client 356
  startup procedures z/OS NFS server 353
saving MVS files 28
saving of mount points 53, 70
sdump macro 254
sdump service 238
sec parameter
  in the exports data set 136
secure(udp) client attribute 83
security 124
security and z/OS UNIX System Services 38
security attribute 51, 59, 175
security authorization facility (SAF) 217
Security Authorization Facility (SAF) 175
security exit 184
security keyword 252
security options, none, saf, safexp, exports 115
security site attribute 115
security, user specified port range support 12
security(exports) attribute 126
security(none) attribute 126
security(saf) 125
security(saf) attribute 361
security(safexp) 125
security(safexp) attribute 361
sequential access method striped data sets 202
sequential byte stream 28
server attributes 87
server log data set, setting up 379
server messages 241
server processes, start number 111
server, bypasses saf checking 104
server, HFS vnode token 106
server, return first character '/' 107
server, start/no start PCNFSD 112
server, z/OS NFS
  control files 7
  creating conventional MVS data sets 6
server, z/OS NFS introduction 3
server/client, user specified port range support 12
service level keyword 229
set the number of times, retransmit 82
setattr request 34
setownernobody processing attribute 97
setownerroot processing attribute 97
setting time stamps 320
setup files, user specified port range support 12
sfmax site attribute 115
share options, cross-region and cross-system 91
shareoptions 6
shareoptions data set attribute 91
shell environment 59
shortcut keys 395
showattr 119
showattr client command, installing 147
showattr command 16, 17, 37, 48, 50, 51, 54, 55, 60, 64, 76, 276, 277
showmount command 62, 75
shutdown of the client 158
shutdown of the server 159
side file 30
side file data set 274
side files, forces 167
side files, mapfile 273
sidefile processing attribute 97
sidefile, allocated storage 115
sidefile(dsname) attribute 30
site attribute
  attrtimeout 108, 109
  bufhigh 104, 108
  cachewindow 104
  checklist 104
  denyrw 105
  dhcp 105
  fileidsize 105
  fn_delimiter 105
  hfs 106
  hfsfbtimeout 106
  hfssec 107
  leadswitch 107
  leasetime 107
  logicalcache 104, 108
  logout 108
  maxrdforszleft 108
  maxtimeout 108, 109

site attribute *(continued)*
mintasks 109
mintimeout 108, 109
mixcase 109
mvssec 110
nfstasks 111
nlm 112
noattrtimeout 108
nochecklist 104
nodenyrw 105
nodhcp 105
noleadswitch 107
nomaxtimeout 108
nopcnfsd 112
nopublic site 113
noreadtimeout 108
norec878 112
nowritetimeout 108
pcnfsd 112
percentsteal 104, 112
public site 113
pubsec 113
readaheadmax 114
readdirtimeout 114
readtimeout 109
restimeout 114
security 115
sfmax 115
smf 116
upcase 109
writetimeout 108, 109
site attributes 104, 117, 214
size, file size determination 313
SMF (System Management Facilities) 130
smf operand 172
smf site attribute 116
smf=on 260
smfwtm macro 131
SMS_keyword 251
soft|hard client attribute 83
Solaris, downloading client commands 150
source code, downloading to the client 359
space data set attribute 91
space, unused 91
spanned data set attribute 91
spanned records, control intervals 91
specifying
Network File System attributes 133
the exports data set 134
srv_ccsid client attribute 83
srv_ccsid message 286
srv_ccsid parameter 11, 66, 67, 83, 127
srv_ccsid processing attribute 98, 100, 101
srv_ccsid setting 67
stable parameter 37
stacks, TCP/IP 130
stale NFS file handle 169
start command 135, 159
start number, server processes 111
start PCNFSD server 112
starting an MVS session 64
starting the client 157
starting the NFS NSM, NLM 161
starting the server 129, 158
statd protocol 327
state 25
statelessness 24
statfs request 34
statistics, SMF 130
status operand 168, 260, 270
STEPLIB DD statement 242
stop command 159
stop mvsnfsc command 158
stop operand 160
stopping subtasks 252
stopping the client 158
stopping the server 160
storage block 174, 184
storage block sample, GFSAUDSA 173
storage considerations 217
storage limit, client 80
storclas data set attribute 91
STOW macro 262, 263
striped files, SAM 21
subtype 7 260
subtype 7, record type-42 131
subtype 8 260
subtype 8, record type-42 131
Sun Solaris–retrieve commands 150
suspending mount processing 166
switchlog 274
symbolic link (symlink) 34
symbolic links 10
SYNAD error, message GFSA877I 268
sync attribtue 36
sync attribute 214
sync processing attribute 98
synchronous write to z/OS UNIX file 37
syntax diagrams, how to read xvi
syntax summary (AIX, UNIX) 47
syntax summary, z/OS command 59
sys parameter
in the exports data set 137
on hfssec site attribute 107
on mvssec site attribute 110
on pubsec site attribute 113
SYSERR DD statement 379
sysplex
multiple NFS servers in 160
system authentication (authsys) security 123
System Authorization Facility (SAF) 8, 123, 124, 125, 138, 261
System Display and Search Facility (SDSF) 206
System Management Facilities (SMF) 130
system toolkits 154, 155
system-managed facility (SMF) 260

T

tag option 286
tag processing attribute 98
tar utility 149
task control block (TCB) 252
task input/output table (TIOT) 295
tbytes value 321
tcp parameter, proto client attribute 82
TCP/IP tuning 202
TCP/IP, stacks 130
tcpip.ETC.RPC, modifying 142
tcpip.profile and /etc/services files, server setup 12
tcprcvbrfsize 202
tcprecvbfrsize 201
tcpsendbfrsize 202
terse option, showattr command 55, 77
tested clients 9
text mode, trailing blanks 93
text processing attribute 93
text processing mode 28
text processing, z/OS UNIX System Services 39
tfiles value 321
time stamps 317
time_delta value 321
time, data set released after a write 98
timed out, data set 114
timeo client attribute 83
timeo value 213
timeout 53
timeout attributes 99
timeout option 197
timeout, control 106
timeout, maximum 108
timeout, minimum 109
timing out, PDS, PDSE 19
toolkits
development 155
system 154, 155
toolkits, development 153
touch (AIX or UNIX) command 151
TRACE CT command 161
trace, collecting messages for 170
trailing blanks 93
translation table, customized 99
translation table, for z/OS UNIX 141
translation table, for z/OS UNIX System Services 39
translation tables 140
Transmission control Protocol/Internet Protocol (TCP/IP) 278
Transmission Control Protocol/Internet Protocol (TCP/IP) 3, 159, 238, 278, 282
transport protocol, client attribute 83
trks data set attribute 88
trusted attribute 119
TSO ALLOCATE command 67
TSO HELP MOUNT command 59
TSO HELP UNMOUNT command 59
TSO MOUNT command 66, 127
TSO UNMOUNT 262

U

UDP checksum processing 200
udp parameter, proto client attribute 82
udprcvbrfsize 202
udprecvbfrsize 201
udpsendbfrsize 202
UID to user name, mapping 362
umount client command 134
umount command 49
unexpected return code, SMF 260
Unicode 85
Unicode Conversion Services 101

Unicode Services 66, 126, 127, 254, 290
unit data set attribute 92
unit_name parameter, unit data set attribute 92
UNIX client, Network File System 47
UNIX commands 49
  mkdir 17
  vi 52
UNIX data sets, commas 105
UNIX kernel 290
UNIX permission bits, checking 115, 124
UNIX System Services 261, 278, 282, 283, 286
UNIX-style credentials 126
UNMNT 275
unmntall operand 169
unmnthfs operand 169
unmntmvs operand 169
unmount (immediate, force) 72
unmount all mount points 169
unmount command 59, 71, 271, 284
unmount HFS mount points 169
unmount legacy mount points 169
unmount operand 169, 271
unstable parameter 37
upcase site attribute 109
updating your MVS password 63
updating your z/OS password 50
upper case, message display 109
USA ASCII format (CCSID 437) 66, 127
usage statistics, collecting 130
user authentication, procedure 13 363
user datagram protocol (UDP) 257, 292
user exit block (UXB) 175
User Exit Block (UXB) 178
user group (UIG) 51
user ID, file attributes 97
user number (UID) 24, 38, 41, 51
user parameter, smf site attribute 116
user specified port range, set up 143
user storage block, login exit 174
user-specified port range support 12
userfile parameter, smf site attribute 116
UTF-8 data format 11, 39

V

validation of MVS passwords and user IDs 124
Variable Recording Area (VRA) 237
verify operand 272
vers attribute 220
vers client attribute 83
vers, mount request 285
vers(3) parameter 285
vers(x) parameter 9
version operand 172
VFS_MOUNT operation 285
VFS_PFSCTL operation 284
VFS_RECOVER operation 284
VFS_STATFS operation 284
VFS_SYNC operation 284
vi command 16, 26
vi editor 18, 19
virtual storage access method (VSAM) 244
virtual storage access method (VSAM) 4, 6, 15, 19, 67, 81, 272, 317, 318
vnode tokens 106
vol attribute 249
vol data set attribute 92
volser parameter, vol data set attribute 92
volume data set attribute 92
VSAM 315
VSAM (virtual storage access method) 67
VSAM data sets, record size 90
VSAM KSDS data sets, keys 89
VTAM 238

W

wait keyword 227
wait parameter, retrieve processing attribute 97
WebNFS support 9
wildcard copy restriction 19
window size, specify 104
write operation 26
write request, nonvolative media 98
write statistic 198
writetimeout 213, 297
writetimeout attribute 250
writetimeout message 249
writetimeout processing attribute 98
writetimeout site attribute 108, 109
writetimeout timeout attribute 99
writing
  fixed-length records 28
  text data files 28
wsize 209, 221
wsize client attribute 83
wtmax value 321, 322
wtmult value 321, 322
wtpref value 321, 322

X

x parameter, cln_ccsid client attribute 80
x parameter, srv_ccsid client attribute 83
X'0F3' ABEND 237
X'806' ABEND 237
X'A03' ABEND 237
X'x13' ABEND 237
X'x22' ABEND 237
X'x37' ABEND 237
X'x3E' ABEND 237
XDR (External Data Representation) 4
xlat attribute 140, 141
xlat client attribute 83
xlat keyword 101
xlat parameter 66, 67, 127
xlat processing attribute 99
xlat() parameter 255
xlat(Y) option 286
xlat=N parameter 67
xlat=Y parameter 67
xreg parameter, shareoptions data set attribute 91
xsys parameter, shareoptions data set attribute 91

Y

Y|N parameter, attrcaching client attribute 80
Y|N parameter, datacaching client attribute 81
Y|N parameter, xlat client attribute 83

Z

z/OS client, Network File System 59
z/OS command, syntax summary 59
z/OS files, showattr command 54
z/OS NFS client startup procedures, sample 356
z/OS NFS server startup procedures, sample 353
z/OS UNIX attribute 36
z/OS UNIX file system 35
z/OS UNIX segment 361
z/OS UNIX System Services 3
  accessing 38
  delete external link 41
  display external link 41
  external link to MVS 40
  mount command 38
  security 38
z/OS, password and user ID 50
zFS 35
zSeries File system 35
zSeries file system (zFS) 211

Readers' Comments — We'd Like to Hear from You

z/OS
Network File System
Guide and Reference

Publication No. SC26-7417-05

Overall, how satisfied are you with the information in this book?

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Overall satisfaction	☐	☐	☐	☐	☐

How satisfied are you that the information in this book is:

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Accurate	☐	☐	☐	☐	☐
Complete	☐	☐	☐	☐	☐
Easy to find	☐	☐	☐	☐	☐
Easy to understand	☐	☐	☐	☐	☐
Well organized	☐	☐	☐	☐	☐
Applicable to your tasks	☐	☐	☐	☐	☐

Please tell us how we can improve this book:

Thank you for your responses. May we contact you? ☐ Yes ☐ No

When you send comments to IBM, you grant IBM a nonexclusive right to use or distribute your comments in any way it believes appropriate without incurring any obligation to you.

Name _____ Address _____

Company or Organization _____

Phone No. _____

Readers' Comments — We'd Like to Hear from You
SC26-7417-05

Cut or Fold
Along Line

Fold and Tape | **Please do not staple** | Fold and Tape

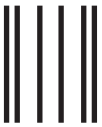

NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES

BUSINESS REPLY MAIL

FIRST-CLASS MAIL PERMIT NO. 40 ARMONK, NEW YORK

POSTAGE WILL BE PAID BY ADDRESSEE

International Business Machines Corporation
Department 55JA, Mail Station P384
2455 South Road
Poughkeepsie, NY 12601-5400

Fold and Tape | **Please do not staple** | Fold and Tape

Cut or Fold
Along Line

Program Number: 5694-A01

Printed in USA

SC26-7417-05