

Managing Network File Systems in Oracle® Solaris 11.3

ORACLE®

Part No: E54789
September 2018

Part No: E54789

Copyright © 2002, 2018, Oracle and/or its affiliates. All rights reserved.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and documentation may provide access to or information about content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services unless otherwise set forth in an applicable agreement between you and Oracle. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services, except as set forth in an applicable agreement between you and Oracle.

Access to Oracle Support

Oracle customers that have purchased support have access to electronic support through My Oracle Support. For information, visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info> or visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs> if you are hearing impaired.

Référence: E54789

Copyright © 2002, 2018, Oracle et/ou ses affiliés. Tous droits réservés.

Ce logiciel et la documentation qui l'accompagne sont protégés par les lois sur la propriété intellectuelle. Ils sont concédés sous licence et soumis à des restrictions d'utilisation et de divulgation. Sauf stipulation expresse de votre contrat de licence ou de la loi, vous ne pouvez pas copier, reproduire, traduire, diffuser, modifier, accorder de licence, transmettre, distribuer, exposer, exécuter, publier ou afficher le logiciel, même partiellement, sous quelque forme et par quelque procédé que ce soit. Par ailleurs, il est interdit de procéder à toute ingénierie inverse du logiciel, de le désassembler ou de le décompiler, excepté à des fins d'interopérabilité avec des logiciels tiers ou tel que prescrit par la loi.

Les informations fournies dans ce document sont susceptibles de modification sans préavis. Par ailleurs, Oracle Corporation ne garantit pas qu'elles soient exemptes d'erreurs et vous invite, le cas échéant, à lui en faire part par écrit.

Si ce logiciel, ou la documentation qui l'accompagne, est livré sous licence au Gouvernement des Etats-Unis, ou à quiconque qui aurait souscrit la licence de ce logiciel pour le compte du Gouvernement des Etats-Unis, la notice suivante s'applique :

U.S. GOVERNMENT END USERS: Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

Ce logiciel ou matériel a été développé pour un usage général dans le cadre d'applications de gestion des informations. Ce logiciel ou matériel n'est pas conçu ni n'est destiné à être utilisé dans des applications à risque, notamment dans des applications pouvant causer un risque de dommages corporels. Si vous utilisez ce logiciel ou ce matériel dans le cadre d'applications dangereuses, il est de votre responsabilité de prendre toutes les mesures de secours, de sauvegarde, de redondance et autres mesures nécessaires à son utilisation dans des conditions optimales de sécurité. Oracle Corporation et ses affiliés déclinent toute responsabilité quant aux dommages causés par l'utilisation de ce logiciel ou matériel pour des applications dangereuses.

Oracle et Java sont des marques déposées d'Oracle Corporation et/ou de ses affiliés. Tout autre nom mentionné peut correspondre à des marques appartenant à d'autres propriétaires qu'Oracle.

Intel et Intel Xeon sont des marques ou des marques déposées d'Intel Corporation. Toutes les marques SPARC sont utilisées sous licence et sont des marques ou des marques déposées de SPARC International, Inc. AMD, Opteron, le logo AMD et le logo AMD Opteron sont des marques ou des marques déposées d'Advanced Micro Devices. UNIX est une marque déposée de The Open Group.

Ce logiciel ou matériel et la documentation qui l'accompagne peuvent fournir des informations ou des liens donnant accès à des contenus, des produits et des services émanant de tiers. Oracle Corporation et ses affiliés déclinent toute responsabilité ou garantie expresse quant aux contenus, produits ou services émanant de tiers, sauf mention contraire stipulée dans un contrat entre vous et Oracle. En aucun cas, Oracle Corporation et ses affiliés ne sauraient être tenus pour responsables des pertes subies, des coûts occasionnés ou des dommages causés par l'accès à des contenus, produits ou services tiers, ou à leur utilisation, sauf mention contraire stipulée dans un contrat entre vous et Oracle.

Accès aux services de support Oracle

Les clients Oracle qui ont souscrit un contrat de support ont accès au support électronique via My Oracle Support. Pour plus d'informations, visitez le site <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info> ou le site <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs> si vous êtes malentendant.

Contents

| | |
|--|----|
| Using This Documentation | 11 |
| 1 About Network File Systems | 13 |
| About the NFS Service | 13 |
| About Autofs | 14 |
| NFS Terminology | 15 |
| NFS Servers and Clients | 15 |
| NFS File Systems | 16 |
| Features of the NFS Service | 16 |
| NFS Version 2 Protocol | 16 |
| NFS Version 3 Protocol | 16 |
| NFS Version 4 Protocol | 17 |
| Controlling NFS Versions | 18 |
| NFS ACL Support | 18 |
| NFS Over TCP | 19 |
| NFS Over UDP | 19 |
| Overview of NFS Over RDMA | 19 |
| Network Lock Manager and NFS | 19 |
| NFS Large File Support | 20 |
| NFS Client Failover | 20 |
| Kerberos Support for the NFS Service | 20 |
| WebNFS Support | 20 |
| RPCSEC_GSS Security Flavor | 21 |
| Extensions for NFS Mounting | 21 |
| Security Negotiation for the WebNFS Service | 21 |
| NFS Server Logging | 22 |
| Autofs Features | 22 |
| Significant Changes in the Oracle Solaris 11 Release | 22 |

| | |
|--|----|
| 2 Network File System Features | 25 |
| How the NFS Service Works | 25 |
| NFS Over RDMA | 26 |
| Version Negotiation in NFS | 28 |
| Features in NFS Version 4 | 28 |
| UDP and TCP Negotiation | 38 |
| File Transfer Size Negotiation | 39 |
| How File Systems Are Mounted in NFS Version 3 | 39 |
| Effects of the -public Option and NFS URLs When Mounting | 40 |
| Client-Side Failover | 41 |
| How NFS Server Logging Works | 43 |
| How the WebNFS Service Works | 43 |
| How WebNFS Security Negotiation Works | 44 |
| WebNFS Limitations With Web Browser Use | 45 |
| Secure NFS Systems | 45 |
| How Mirror Mounts Work | 49 |
| Mounting a File System Using Mirror Mounts | 49 |
| Unmounting a File System Using Mirror Mounts | 49 |
| How NFS Referrals Work | 50 |
| When to Use NFS Referrals | 50 |
| Creating an NFS Referral | 50 |
| Removing an NFS Referral | 51 |
| How Autofs Works | 51 |
| How Autofs Navigates Through the Network | 53 |
| Autofs Maps | 53 |
| How Autofs Starts the Navigation Process | 59 |
| Autofs Mount Process | 60 |
| How Autofs Selects the Nearest Read-Only Files for Clients | 61 |
| Autofs and Weighting | 65 |
| Variables in an Autofs Map Entry | 65 |
| Maps That Refer to Other Maps | 66 |
| Executable Autofs Maps | 67 |
| Default Autofs Behavior With Name Services | 67 |
| Autofs Reference | 69 |
| Autofs and Metacharacters | 69 |
| Autofs and Special Characters | 70 |

| | |
|--|----|
| 3 Administering Network File Systems | 71 |
| About Administering Network File Systems | 71 |
| Automatic File System Sharing | 72 |
| Task Map for File System Sharing | 72 |
| ▼ How to Set Up Automatic File System Sharing | 72 |
| ▼ How to Enable NFS Server Logging | 73 |
| Mounting File Systems | 74 |
| Mounting File Systems Task Map | 74 |
| ▼ How to Mount a File System at Boot Time | 75 |
| ▼ How to Mount a File System From the Command Line | 76 |
| Mounting With the Automounter | 77 |
| ▼ How to Mount All File Systems From a Server | 77 |
| ▼ How to Use Client-Side Failover | 78 |
| ▼ How to Disable Mount Access for One Client | 78 |
| ▼ How to Mount an NFS File System Through a Firewall | 79 |
| Mount an NFS File System by Using an NFS URL | 79 |
| Displaying Information About File Systems Available for Mounting | 80 |
| Setting Up the NFS Service | 81 |
| Starting and Stopping the NFS Service | 81 |
| Starting and Stopping the Automounter | 81 |
| Selecting Different Versions of NFS | 82 |
| Administering the Secure NFS System | 85 |
| ▼ How to Set Up a Secure NFS Environment With DH Authentication | 85 |
| Administering WebNFS | 87 |
| Planning for WebNFS Access | 87 |
| ▼ How to Enable WebNFS Access | 88 |
| Accessing an NFS URL by Using a Browser | 89 |
| Enabling WebNFS Access Through a Firewall | 90 |
| Administering NFS Referrals | 90 |
| ▼ How to Create and Access an NFS Referral | 90 |
| ▼ How to Remove an NFS Referral | 91 |
| Administering FedFS | 91 |
| Set Up a DNS Record for a FedFS Server | 91 |
| ▼ How to Create a Namespace Database | 92 |
| ▼ How to Use a Secured Connection to the NSDB | 93 |
| ▼ How to Create a FedFS Referral | 94 |

| | |
|--|-----|
| 4 Administering Autofs | 95 |
| Autofs Administration | 95 |
| Using SMF Parameters to Configure Your Autofs Environment | 97 |
| ▼ How to Configure Your Autofs Environment Using SMF Parameters | 97 |
| Administrative Tasks Involving Maps | 97 |
| Modifying the Maps | 98 |
| ▼ How to Modify the Master Map | 98 |
| ▼ How to Modify Indirect Maps | 99 |
| ▼ How to Modify Direct Maps | 99 |
| Avoiding Mount Point Conflicts | 99 |
| Accessing Non-NFS File Systems | 100 |
| Customizing the Automounter | 100 |
| Setting Up a Common View of /home | 100 |
| ▼ How to Set Up /home With Multiple Home Directory File Systems | 101 |
| ▼ How to Consolidate Project-Related Files Under a Common Directory | 102 |
| ▼ How to Set Up Different Architectures to Access a Shared Namespace | 104 |
| ▼ How to Support Incompatible Client Operating System Versions | 105 |
| ▼ How to Replicate Shared Files Across Several Servers | 106 |
| Autofs Security Restrictions | 106 |
| ▼ How to Use a Public File Handle With Autofs | 107 |
| ▼ How to Use NFS URLs With Autofs | 107 |
| Disabling Autofs Browsability | 107 |
| | |
| 5 Commands for Managing Network File Systems | 111 |
| NFS Commands | 111 |
| automount Command | 112 |
| clear_locks Command | 113 |
| fsstat Command | 113 |
| mount Command | 114 |
| umount Command | 120 |
| mountall Command | 121 |
| umountall Command | 121 |
| sharectl Command | 122 |
| share Command | 124 |
| unshare Command | 129 |
| shareall Command | 130 |

| | |
|--|------------|
| unshareall Command | 130 |
| showmount Command | 130 |
| nfsref Command | 132 |
| FedFS Commands | 132 |
| 6 Troubleshooting Network File Systems | 133 |
| Strategies for NFS Troubleshooting | 133 |
| Commands for Troubleshooting NFS Problems | 134 |
| nfsstat Command | 134 |
| pstack Command | 137 |
| rpcinfo Command | 137 |
| snoop Command | 139 |
| truss Command | 140 |
| NFS Troubleshooting Procedures | 140 |
| ▼ How to Check Connectivity on an NFS Client | 141 |
| ▼ How to Check the NFS Server Remotely | 142 |
| ▼ How to Verify the NFS Service on the Server | 143 |
| ▼ How to Restart NFS Service | 145 |
| Identifying the Host Providing the NFS Service | 145 |
| ▼ How to Verify mount Command Options | 145 |
| Troubleshooting Autofs | 146 |
| Error Messages Generated by automount -v | 146 |
| Miscellaneous Error Messages | 148 |
| Other Errors With Autofs | 150 |
| NFS Error Messages | 150 |
| 7 Accessing Network File Systems | 157 |
| NFS Files | 157 |
| /etc/default/nfslogd File | 158 |
| /etc/nfs/nfslog.conf File | 159 |
| NFS Daemons | 160 |
| automountd Daemon | 161 |
| lockd Daemon | 162 |
| mountd Daemon | 163 |
| nfs4cbd Daemon | 164 |
| nfsd Daemon | 164 |

| | |
|---|------------|
| nfslogd Daemon | 165 |
| nfsmapid Daemon | 165 |
| rearsed Daemon | 172 |
| statd Daemon | 172 |
| A NFS File Sharing Command Reference | 173 |
| Administering Network File Systems | 173 |
| Troubleshooting Network File Systems | 174 |
| Index | 175 |

Using This Documentation

- **Overview** – Describes how to administer and access the network file systems.
- **Audience** – Technicians, system administrators, and authorized service providers.
- **Required knowledge** – Basic and some advanced network administration skills.

Product Documentation Library

Documentation and resources for this product and related products are available at <http://www.oracle.com/pls/topic/lookup?ctx=E53394-01>.

Feedback

Provide feedback about this documentation at <http://www.oracle.com/goto/docfeedback>.

◆◆◆ CHAPTER 1

About Network File Systems

This chapter provides an overview of the network file system (NFS) service, which can be used to access file systems over the network. The NFS service enables any system to access any other systems' file systems. A system can assume the role of client, server, or both, at any particular time on a network. Autofs is a client-side service used to mount the file systems that are shared through the NFS service. Autofs is a file system structure that provides automatic mounting. The chapter includes a discussion of the concepts necessary to understand the NFS service and a description of the latest features in NFS and autofs.

This chapter contains the following topics:

- [“About the NFS Service” on page 13](#)
- [“About Autofs” on page 14](#)
- [“NFS Terminology” on page 15](#)
- [“Features of the NFS Service” on page 16](#)
- [“Significant Changes in the Oracle Solaris 11 Release” on page 22](#)

Note - If your system has zones enabled and you want to use NFS in a non-global zone, see [“Non-Global Zones as NFS Clients” in *Creating and Using Oracle Solaris 10 Zones*](#).

About the NFS Service

The NFS service enables systems of different architectures that run different operating systems to share file systems across a network.

The NFS environment can be implemented on different operating systems because NFS defines an abstract model of a file system rather than an architectural specification. Each OS applies the NFS model to its file system semantics. This model means that file system operations, such as reading and writing, function as though the operations are accessing a local file.

The NFS service has the following benefits:

- Enables multiple systems to use the same files so that everyone on the network can access the same data
- Reduces storage costs by having systems share applications instead of requiring local disk space for each user application
- Provides data consistency and reliability because all users can read the same set of files
- Makes mounting of file systems transparent to users
- Makes accessing of remote files transparent to users
- Supports heterogeneous environments
- Reduces system administration overhead

The NFS service makes the physical location of a file system irrelevant to the user. Instead of having to place commonly used files on every system, you can share the original file from the NFS server's file system. All other systems access the files across the network. Under NFS operation, remote file systems are almost indistinguishable from local file systems.

About Autofs

File systems that are shared through the NFS service can be mounted by using automatic mounting. Autofs, a client-side service, is a file system structure that provides automatic mounting. The autofs file system is initialized by automount, which is run automatically when a system is booted. The automount daemon, automountd, runs continuously, mounting and unmounting remote file systems as necessary.

Whenever a client system that is running automountd tries to access a remote file system, the daemon mounts the remote file system. This remote file system remains mounted for as long as needed. If the remote file system is not accessed for a certain period of time, it is automatically unmounted.

Mounting is not required at boot time, and the user no longer has to know the superuser password to mount a directory. Users do not need to use the mount and umount commands. The autofs service mounts and unmounts file systems as required without any intervention by the user.

Mounting some file systems with the automountd command does not exclude the possibility of mounting other systems with the mount command. A diskless computer *must* mount / (root), /usr, and /usr/kvm through the mount command and the /etc/vfstab file.

For more information about the autofs service, see:

- [Table 6, “Tasks for Administering Autofs,” on page 95](#)
- [“How Autofs Works” on page 51](#)

NFS Terminology

This section includes basic terminology that must be understood to work with the NFS service. For more information about the NFS service, see [Chapter 5, “Commands for Managing Network File Systems”](#).

IP addresses that are used in this guide conform to [RFC 5737, IPv4 Address Blocks Reserved for Documentation](#) and [RFC 3849, IPv6 Address Prefix Reserved for Documentation](#).

- IPv4 addresses used in this documentation are blocks 192.0.2.0/24, 198.51.100.0/24, and 203.0.113.0/24.
To show a subnet, the block is divided into multiple subnets by borrowing enough bits from the host to create the required subnet. For example, host address 192.0.2.0 might have subnets 192.0.2.32/27 and 192.0.2.64/27.
- IPv6 addresses have prefix 2001:DB8::/32.

NFS Servers and Clients

The terms *client* and *server* describe the roles that a system assumes when sharing file systems. Servers are systems that share their file systems over a network. The systems that access these file systems are the clients.

Clients access files on the server by mounting the server's shared file systems. When a client mounts a remote file system, the client does not make a copy of the file system. Rather, the mounting process uses a series of remote procedure calls that enable the client to access the server's shared file system transparently. The mount resembles a local mount. You can type commands as if the file systems are local. For more information, see [“Mounting File Systems” on page 74](#).

After a file system has been shared on a server through an NFS operation, the file system can be accessed from a client. You can mount an NFS file system automatically with `autofs`. For information about `autofs`, see [“About Autofs” on page 14](#). For information about tasks that involve the `share` command and `autofs`, see [“Automatic File System Sharing” on page 72](#) and [Table 6, “Tasks for Administering Autofs,” on page 95](#).

NFS File Systems

The objects that can be shared with the NFS service include any whole or partial directory tree or a file hierarchy, including a single file. A system cannot share a file hierarchy that overlaps a file hierarchy that is already shared. Peripheral devices such as modems and printers cannot be shared.

In most UNIX system environments, a file hierarchy that can be shared corresponds to a file system or to a portion of a file system. However, because NFS works across operating systems and the concept of a file system might be meaningless in non-UNIX environments, the term *file system* refers to a file or file hierarchy that can be shared and mounted with NFS.

Features of the NFS Service

This section describes important features of the NFS service.

NFS Version 2 Protocol

NFS Version 2, the first version of the NFS protocol is widely used. All Oracle Solaris releases support the NFS Version 2 protocol.

NFS Version 3 Protocol

Unlike the NFS Version 2 protocol, the NFS Version 3 protocol can handle files that are larger than 2 GB. For information about handling large files in NFS, see [“NFS Large File Support” on page 20](#).

The NFS Version 3 protocol enables safe asynchronous writes on the server, which improves performance by allowing the server to cache client write requests in memory. The client no longer waits for the server to commit the changes to disk, so the response time is faster. Also, the server can batch the requests, which improves the response time on the server.

Many Oracle Solaris NFS Version 3 operations return the file attributes, which are stored in the local cache. Because the cache is updated more often, the requirement to perform a separate operation to update this data arises less often. Therefore, the number of Remote Procedure Calls (RPC) to the server is reduced, improving performance.

The process for verifying file access permissions has been improved. Version 2 generated a "write error" message or a "read error" message if users tried to copy a remote file without the

appropriate permissions. In Version 3, the permissions are checked before the file is opened, so the error is reported as an "open error".

The NFS Version 3 protocol removes the 8 KB transfer size limit. Clients and servers can negotiate whatever transfer size the clients and servers support, rather than conform to the 8 KB limit that Version 2 imposed. Note that in earlier Oracle Solaris implementations, the protocol defaulted to a 32 KB transfer size. Starting with the Oracle Solaris 10 release, restrictions on wire transfer sizes were relaxed. The transfer size is based on the capabilities of the underlying transport.

NFS Version 4 Protocol

The NFS Version 4 protocol represents the user ID and the group ID as strings. The `nfsmapid` daemon is used by the NFS Version 4 client and server for the following mappings:

- Map the user ID and group ID strings to local numeric IDs
- Map the local numeric IDs to user ID and group ID strings

For more information about the `nfsmapid` daemon, see [“NFS Daemons” on page 160](#).

Note that in NFS Version 4, the `nfsmapid` daemon, is used to map user IDs or group IDs in Access Control List (ACL) entries on a server to user IDs or group IDs in ACL entries on a client. The reverse is also true. For more information about user ID and group ID mapping, see [“ACLs and `nfsmapid` in NFS Version 4” on page 37](#) and [“NFS ACL Support” on page 18](#).

With NFS Version 4, when you unshare a file system, all the state information for any open files or file locks in that file system is destroyed. In NFS Version 3, the server maintains any locks that the clients had obtained before the file system was unshared. For more information about unsharing a file system in NFS Version 4, see [“Unsharing and Resharing a File System in NFS Version 4” on page 29](#).

NFS Version 4 servers use a pseudo file system to provide clients with access to exported objects on the server. For more information about pseudo file system, see [“File System Namespace in NFS Version 4” on page 29](#). NFS Version 4 supports volatile file handles. For more information, see [“Volatile File Handles in NFS Version 4” on page 32](#).

Delegation, a technique by which the server delegates the management of a file to a client, is supported on both the client and the server. For example, the server could grant either a read delegation or a write delegation to a client. For more information about delegation, see [“Delegation in NFS Version 4” on page 35](#).

NFS Version 4 does not support LIPKEY/SPKM security.

Also, NFS Version 4 does not use the following daemons:

- lockd
- nfslogd
- statd

For a complete list of the features in NFS Version 4, see [“Features in NFS Version 4” on page 28](#).

For information about setting up the NFS services, see [“Setting Up the NFS Service” on page 81](#).

Controlling NFS Versions

The SMF repository includes parameters to control the NFS protocols that are used by both the client and the server. For example, you can use parameters to manage version negotiation. For more information about the client and server parameters, see [“NFS Daemons” on page 160](#). For more information about the parameter values for NFS daemons, see the [nfs\(4\)](#) man page.

NFS ACL Support

Access control list (ACL) provides a mechanism to set file access permissions instead of using the standard UNIX file permissions. NFS ACL support provides a method of changing and viewing ACL entries from an Oracle Solaris NFS client to an Oracle Solaris NFS server.

The NFS Version 2 and Version 3 implementations support the old POSIX-draft style ACLs. POSIX-draft ACLs are natively supported by UFS. For more information about POSIX-draft ACLs, see [“Using Access Control Lists to Protect UFS Files” in *Securing Files and Verifying File Integrity in Oracle Solaris 11.3*](#).

The NFS Version 4 protocol supports NFS Version 4 style ACLs. NFS Version 4 ACLs are natively supported by Oracle Solaris ZFS. You must use ZFS as the underlying file system on the NFS Version 4 server for full featured NFS Version 4 ACL functionality. NFS Version 4 ACLs have a rich set of inheritance properties, as well as a set of permission bits beyond the standard read, write, and execute. For more information about using ACLs to protect ZFS files, see [“Setting ACLs on ZFS Files” in *Securing Files and Verifying File Integrity in Oracle Solaris 11.3*](#). For more information about support for ACLs in NFS Version 4, see [“ACLs and nfmapiid in NFS Version 4” on page 37](#).

NFS Over TCP

The default transport protocol for the NFS protocol is TCP (Transmission Control Protocol). TCP helps performance on slow networks and wide area networks. TCP also provides congestion control and error recovery. NFS over TCP works with the NFS Version 2, NFS Version 3, and NFS Version 4 protocols.

Note - If InfiniBand hardware is available on the system, the default transport protocol changes from TCP to the Remote Direct Memory Access (RDMA) protocol. For more information, see [“Overview of NFS Over RDMA” on page 19](#) and [“NFS Over RDMA” on page 26](#). Note that, if you use the `proto=tcp` mount option, NFS mounts are forced to use TCP only.

NFS Over UDP

Starting with the Oracle Solaris 11 release, the NFS client uses only one UDP (User Datagram Protocol) reserved port, which is configurable. The system can be configured to use more than one port to increase system performance. This capability mirrors NFS over TCP support, which has been configurable in this way since its inception. For more information about tuning the NFS environment, see [Oracle Solaris 11.3 Tunable Parameters Reference Manual](#).

Overview of NFS Over RDMA

If InfiniBand hardware is available on the system, the default transport protocol changes from TCP to the RDMA protocol. The RDMA protocol is a technology for memory-to-memory transfer of data over high-speed networks. Specifically, RDMA provides remote data transfer directly to and from memory without CPU intervention. To provide this capability, RDMA combines the interconnect I/O technology of InfiniBand with the Oracle Solaris OS. However, if you use the `proto=tcp` mount option, NFS mounts are forced to use TCP only. For more information about using the RDMA protocol for NFS, see [“NFS Over RDMA” on page 26](#).

Network Lock Manager and NFS

The Network Lock Manager provides UNIX record locking for any files being shared over NFS. The locking mechanism enables clients to synchronize their I/O requests with other clients, ensuring data integrity.

Note - The Network Lock Manager is used only for NFS Version 2 and NFS Version 3 mounts. File locking is built into the NFS Version 4 protocol.

NFS Large File Support

The NFS Version 3 protocol can handle files that are larger than 2 GB, but the NFS Version 2 protocol cannot.

NFS Client Failover

Dynamic failover of read-only file systems provide a high level of availability for read-only resources that are already replicated, such as man pages, other documentation, and shared binaries. Failover can occur any time after the file system is mounted. Manual mounts can now list multiple replicas, much like the automounter in previous releases. The automounter has not changed, except that failover no longer waits until the file system is remounted. For more information, see [“How to Use Client-Side Failover” on page 78](#) and [“Client-Side Failover” on page 41](#).

Kerberos Support for the NFS Service

The NFS service supports Kerberos Version 5 authentication, integrity, and privacy when you configure NFS clients and servers to support Kerberos. You can use the mount and share command-line options when you use Kerberos for secure authentication. For information about Kerberos Version 5 authentication, see [“Configuring Kerberos NFS Servers” in *Managing Kerberos and Other Authentication Services in Oracle Solaris 11.3*](#).

WebNFS Support

WebNFS provides the capability to make a file system on the Internet accessible through firewalls. This capability uses an extension to the NFS protocol. One advantage of using the WebNFS™ protocol for Internet access is its reliability. The service is built as an extension of the NFS Version 3 and Version 2 protocol. Additionally, WebNFS enables you to share files without the administrative overhead of an anonymous ftp site. For more information about WebNFS, see [“Security Negotiation for the WebNFS Service” on page 21](#) and [“Administering WebNFS” on page 87](#).

Note - The NFS Version 4 protocol is preferred over the WebNFS service. NFS version 4 fully integrates all the security negotiation that was added to the MOUNT protocol and the WebNFS service.

RPCSEC_GSS Security Flavor

A security flavor called `RPCSEC_GSS` uses the standard GSS-API interfaces to provide authentication, integrity, and privacy, as well as enabling support of multiple security mechanisms. For more information about support of Kerberos V5 authentication, see [“Kerberos Support for the NFS Service” on page 20](#). For more information about GSS-API, see Chapter 4, [“Writing Applications That Use GSS-API” in *Developer’s Guide to Oracle Solaris 11.3 Security*](#).

Extensions for NFS Mounting

The NFS service provides extensions to the `mount` and `automountd` commands in Oracle Solaris. These extensions enable the `mount` request to use the public file handle instead of the MOUNT protocol. The WebNFS service uses the MOUNT protocol as the access method. By using the public file handle, the `mount` can occur through a firewall. As there are fewer transactions between the server and the client, the `mount` occurs faster.

The extensions also enable NFS URLs to be used instead of the standard path name. Also, you can use the `public` option with the `mount` command and the `automounter` maps to force the use of the public file handle. For more information about the WebNFS service, see [“WebNFS Support” on page 20](#).

Security Negotiation for the WebNFS Service

The NFS service enables a WebNFS client to negotiate a security mechanism with an NFS server. WebNFS client uses a protocol to negotiate a security mechanism with an NFS server. This protocol enables you to use secure transactions with the WebNFS service. For more information about security negotiation for WebNFS, see [“How WebNFS Security Negotiation Works” on page 44](#).

NFS Server Logging

Note - NFS Version 4 does not support the server logging feature.

NFS server logging enables an NFS server to provide a record of file operations that have been performed on its file systems. The record includes information about which file was accessed, when the file was accessed, and who accessed the file. You can specify the location of the logs that contain this information through a set of configuration options. You can also use these options to select the operations to be logged. The NFS server logging feature is particularly useful for sites that make anonymous FTP archives available to NFS and WebNFS clients. For more information, see [“How to Enable NFS Server Logging” on page 73](#).

Autofs Features

Autofs works with file systems that are specified in the local namespace. This information can be maintained in LDAP, NIS (Network Information Service), or local files. Autofs supports the following features:

- A fully multithreaded version of the automountd feature capability makes autofs reliable. This feature enables concurrent servicing of multiple mounts, which prevents the service from hanging if a server is unavailable.
- The automountd feature also provides on-demand mounting. Only the top file system is mounted. Other file systems that are related to this mount point are mounted when needed.
- The autofs service supports the “browsability” of indirect maps. This support enables a user to see which directories could be mounted without having to actually mount each file system. A -nobrowse option ensures that large file systems, such as /net and /home, are not automatically browsable. Also, you can turn off autofs browsability on each client by using the -n option with the automount command. For more information about different methods to disable autofs browsability, see [“Disabling Autofs Browsability” on page 107](#).

Significant Changes in the Oracle Solaris 11 Release

The Oracle Solaris 11 release includes the following enhancements:

- A new property, `nfs_props/showmount_info`, has been added to the `/network/nfs/server:default` service. This property controls the amount of information that the `showmount` command displays to remote clients. For more information about the `nfs_props/showmount_info` property, see the [`showmount\(1M\)` man page](#).

- Support for the Federated File System (FedFS) referrals has been added. This feature enables referral information for several servers to be centralized in LDAP. For more information about FedFS referrals, see [“Administering FedFS” on page 91](#).
- The configuration properties that used to be set by editing the `/etc/default/autofs` and `/etc/default/nfs` files can now be set in the Service Management Facility (SMF) repository. For more information about the new SMF properties and the daemons using the new SMF properties, see [“NFS Daemons” on page 160](#).
- The NFS service provides support for mirror mounts. Mirror mounts enable an NFS Version 4 client to traverse shared file system mount points in the server namespace. For NFS Version 4 mounts, the automounter performs a mount of the server namespace root and relies on mirror mounts to access its file systems. The main advantage that mirror mounts offer over the automounter is that mounting a file system using mirror mounts does not require the overhead associated with administering automount maps. Mirror mounts provide the following features:
 - Namespace changes are immediately visible to all clients.
 - New shared file systems are discovered instantly and mounted automatically.
 - File systems unmount automatically after a designated inactivity period.

For more information about mirror mounts, see:

- [“How to Mount All File Systems From a Server” on page 77](#)
- [“How Mirror Mounts Work” on page 49](#)
- NFS referrals have been added to the NFS service. Referrals are server-based redirections that an NFS Version 4 client can follow to find a file system. The NFS server supports referrals created by the `nfsref` command. The NFS Version 4 client follows these referrals to mount the file system from the actual location. The creation of referrals replaces the editing of the automounter map. NFS referrals provide these features:
 - All of the features of mirror mounts
 - Automounter-like functionality without any dependence on the automounter
 - No setup required on either the client or server

For more information about NFS referrals, see:

- [“Administering NFS Referrals” on page 90](#)
- [“How NFS Referrals Work” on page 50](#)
- `nfsref(1M)` man page
- The ability to mount the per-DNS-domain root of a FedFS namespace has been added. This mount point can be used with NFS referrals to bridge from one file server to another file server, building an arbitrarily large namespace. For more information about FedFS domain root, see:
 - [“Set Up a DNS Record for a FedFS Server” on page 91](#)
 - [“Mount Point /nfs4” on page 55](#)

- The `sharectl` utility enables you to configure and manage file sharing protocols, such as NFS. For example, this utility enables you to set client and server operational properties, display property values for a specific protocol, and obtain the status of a protocol. For more information, see the [sharectl\(1M\)](#) man page.
- The ability to restrict NFS service to clients which use reserved ports, that is port numbers below 1024, has been added. This affects file systems shared with `AUTH_SYS`. This can be set for all shares by using the `resvport` setting with the `sharectl` command, or for individual file systems by using the `resvport` share option.

Network File System Features

This chapter describes the relationship of Remote Direct Memory Access (RDMA) protocol to other transport protocols. RDMA is the default transport for NFS. This chapter also describes how the NFS service works, which includes version negotiation and the features introduced in NFS Version 4 for file sharing.

This chapter contains the following topics:

- “How the NFS Service Works” on page 25
- “How Mirror Mounts Work” on page 49
- “How NFS Referrals Work” on page 50
- “How Autofs Works” on page 51, including “Autofs Maps” on page 53
- “Autofs Reference” on page 69

Note - If your system has zones enabled and you want to use this feature in a non-global zone, see [Introduction to Oracle Solaris Zones](#).

How the NFS Service Works

The following sections describe some of the complex functions of the NFS software. Some of the feature descriptions in this section are exclusive to NFS Version 4.

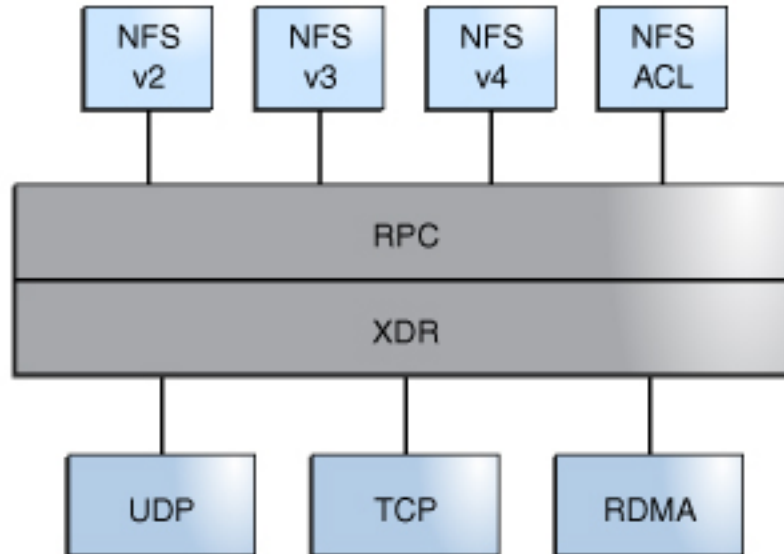
- “NFS Over RDMA” on page 26
- “Version Negotiation in NFS” on page 28
- “Features in NFS Version 4” on page 28
- “UDP and TCP Negotiation” on page 38
- “File Transfer Size Negotiation” on page 39
- “How File Systems Are Mounted in NFS Version 3” on page 39
- “Effects of the `-public` Option and NFS URLs When Mounting” on page 40

- [“Client-Side Failover” on page 41](#)
- [“How NFS Server Logging Works” on page 43](#)
- [“How the WebNFS Service Works” on page 43](#)
- [“How WebNFS Security Negotiation Works” on page 44](#)
- [“WebNFS Limitations With Web Browser Use” on page 45](#)
- [“Secure NFS Systems” on page 45, including “Secure RPC” on page 46](#)

Note - If your system has zones enabled and you want to use this feature in a non-global zone, see [Introduction to Oracle Solaris Zones](#).

NFS Over RDMA

Starting with the Oracle Solaris 11.1 release, the default transport for NFS is the Remote Direct Memory Access (RDMA) protocol. This protocol provides memory-to-memory transfer of data over high-speed networks. Specifically, RDMA provides remote data transfer directly to and from memory without CPU intervention. RDMA also provides direct data placement, which eliminates data copies and therefore further eliminates CPU intervention. Thus, RDMA frees up not only the host CPU, but also reduces contention for the host memory and I/O buses. To provide this capability, RDMA combines the interconnect I/O technology of InfiniBand, which you can use on both SPARC and x86 platforms, with the Oracle Solaris operating system. The following figure shows the relationship of RDMA to other protocols, such as UDP and TCP.

FIGURE 1 Relationship of RDMA to Other Protocols

Because RDMA is the default transport protocol for NFS, no special share or mount options are required to use RDMA on a client or server. The existing automounter maps, `vfstab` and file system shares, work with the RDMA transport. NFS mounts over the RDMA transport occur transparently when InfiniBand connectivity exists between the client and the server. InfiniBand connectivity feature works on both SPARC and x86 platforms. If the RDMA transport is not available on both the client and the server, the TCP transport is the initial fallback, followed by UDP if TCP is unavailable. However, if you use the `proto=rdma` mount option, NFS mounts are forced to use RDMA only.

To specify the use of only TCP and UDP, you can use the `proto=tcp/udp` mount option. This option disables RDMA on an NFS client. For more information about NFS mount options, see the [mount_nfs\(1M\)](#) and [mount\(1M\)](#) man pages.

Note - RDMA for InfiniBand uses the IP addressing format and the IP lookup infrastructure to specify peers. However, because RDMA is a separate protocol stack, it does not fully implement all IP semantics. For example, RDMA does not use IP addressing to communicate with peers. Therefore, RDMA might bypass configurations for various security policies that are based on IP addresses. However, the NFS and RPC administrative policies, such as mount restrictions and secure RPC, are not bypassed.

Version Negotiation in NFS

The NFS initiation process includes negotiating the protocol version levels for servers and clients. If you do not specify the version level, then the best level is selected by default. For example, if both the client and the server can support NFS Version 3, then that is used. If the client or the server can only support NFS Version 2, then that is used.

You can set the `client_versmin`, `client_versmax`, `server_versmin`, and `server_versmax` parameters by using the `sharectl` command. Your specified minimum and maximum values for the server and the client replace the default values for these parameters. For both the client and the server, the default minimum value is 2 and the default maximum value is 4. To find the version supported by the server, the NFS client begins with the value for `client_versmax` and continues to try each version until reaching the version value for `client_versmin`. As soon as the supported version is found, the process terminates. For example, if `client_versmax=4` and `client_versmin=2`, then the client attempts NFS Version 4 first, then NFS Version 3, and finally NFS Version 2. If `client_versmax` and `client_versmin` are set to the same value, then the client always uses this version and does not attempt any other version. If the server does not offer this version, the mount fails.

Note - You can override the values that are determined by the version negotiation in NFS by using the `vers` option with the `mount` command. For more information about available options for the `mount` command, see the [mount_nfs\(1M\)](#) man page.

For information about setting up the NFS service, see [“Setting Up the NFS Service” on page 81](#).

Features in NFS Version 4

This section provides descriptions of the features that were introduced in NFS Version 4:

- [“Unsharing and Resharing a File System in NFS Version 4” on page 29](#)

- [“File System Namespace in NFS Version 4” on page 29](#)
- [“Volatile File Handles in NFS Version 4” on page 32](#)
- [“Client Recovery in NFS Version 4” on page 33](#)
- [“OPEN Share Support in NFS Version 4” on page 34](#)
- [“Delegation in NFS Version 4” on page 35](#)
- [“ACLs and nfsmapid in NFS Version 4” on page 37](#)

Note - Starting with the Oracle Solaris 10 release, NFS Version 4 does not support the LIPKEY/SPKM security flavor. Also, NFS Version 4 does not use the mountd, nfslogd, and statd daemons.

For information about setting up NFS services, see [“Setting Up the NFS Service” on page 81](#).

Unsharing and Resharing a File System in NFS Version 4

With both NFS Version 3 and NFS Version 4, if a client attempts to access a file system that has been unshared, the server responds with an error code. However, with NFS Version 3, the server maintains any locks that the clients had obtained before the file system was unshared. Thus, when the file system is reshared, NFS Version 3 clients can access the file system as though that file system had never been unshared.

With NFS Version 4, when a file system is unshared, all the state information for any open files or file locks in that file system is destroyed. If the client attempts to access these files or locks, the client receives an error. This error is usually reported as an I/O error to the application. However, resharing a currently shared file system to change options does not destroy any of the state information on the server.

For information about client recovery in NFS Version 4, see [“Client Recovery in NFS Version 4” on page 33](#). For information about available options for the unshare command, see the [unshare_nfs\(1M\)](#) man page.

File System Namespace in NFS Version 4

NFS Version 4 servers create and maintain a pseudo file system that provides NFS clients with seamless access to all exported objects on the server. Prior to NFS Version 4, the pseudo file system did not exist. NFS clients were forced to mount each shared server file system for access.

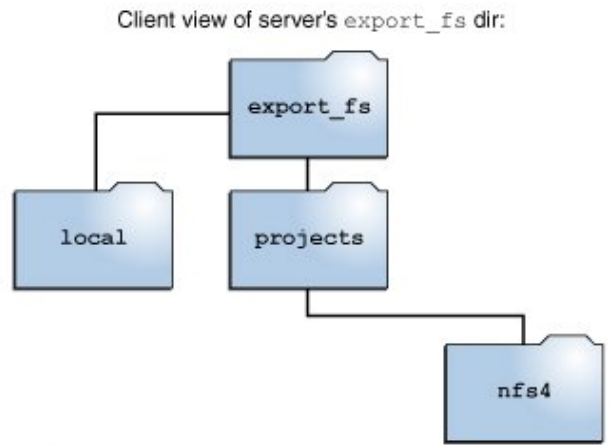
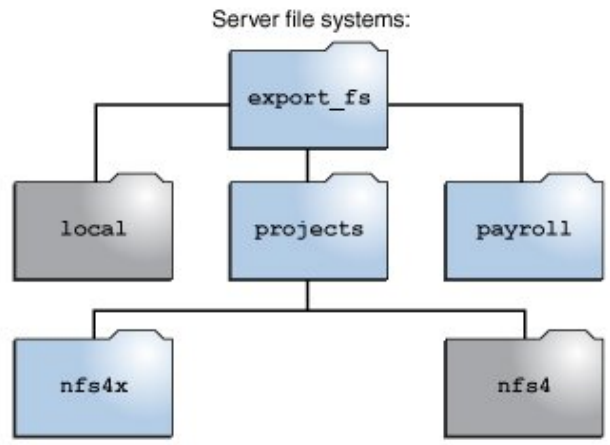
A pseudo file system is a structure that contains only directories and is created by the server. The pseudo file system permits a client to browse the hierarchy of exported file systems. Thus, the client's view of the pseudo file system is limited to paths that lead to exported file systems.

Previous versions of NFS did not permit a client to traverse server file systems without mounting each file system. However, in NFS Version 4, the server namespace does the following:

- Restricts the client's file system view to directories that lead to server exports.
- Provides clients with seamless access to server exports without requiring the client to mount each underlying file system. However, different operating systems might require the client to mount each server file system.

FIGURE 2 Views of the Server File System and the Client File System in NFS Version 4

| | |
|---------------------------------------|-------------------------|
| Server exports: | Server file systems: |
| <code>/export_fs/local</code> | <code>/</code> |
| <code>/export_fs/projects/nfs4</code> | <code>/export_fs</code> |



■ Exported directories

In the example shown in the figure, the client cannot see the `payroll` directory and the `nfs4x` directory because these directories are not exported and do not lead to exported directories. However, the `local` directory is visible to the client because `local` is an exported directory. The `projects` directory is visible to the client because `projects` leads to the exported directory, `nfs4`. Thus, portions of the server namespace that are not explicitly exported are bridged with a pseudo file system that views only the exported directories and those directories that lead to server exports.

Volatile File Handles in NFS Version 4

File handles are created on the NFS server and contain information that uniquely identifies files and directories. In NFS Version 2 and NFS Version 3, the NFS server returned persistent file handles. Thus, the NFS client could guarantee that the NFS server would generate a file handle that always referred to the same file. For example:

- If a file was deleted and replaced with a file of the same name, the server would generate a new file handle for the new file. If the client used the old file handle, the server would return an error that the file handle was stale.
- If a file was renamed, the file handle would remain the same.
- If the server was rebooted, the file handles would remain the same.

Thus, when the server received a request from a client that included a file handle, the resolution was straightforward and the file handle always referred to the correct file.

The method of using persistent file handles for identifying files and directories for NFS operations was fine for most UNIX servers. However, the method could not be implemented on servers that relied on other methods of identification, such as a file's path name. To resolve this problem, the NFS Version 4 protocol permits a server to declare that its file handles are volatile. If the file handle does change, the client must find the new file handle.

Like NFS Versions 2 and NFS Version 3 servers, the Oracle Solaris NFS Version 4 server always provides persistent file handles. However, Oracle Solaris NFS Version 4 clients that access non Oracle Solaris NFS Version 4 servers must support volatile file handles if the server uses them. Specifically, when the server communicates to the client that the file handle is volatile, the client must cache the mapping between the path name and file handle. The client uses the volatile file handle until it expires. On expiration, the client does the following:

- Flushes the cached information that refers to that file handle
- Searches for that file's new file handle
- Retries the operation

Note - The server always communicates to the client which file handles are persistent and which file handles are volatile.

Volatile file handles might expire in any of these situations:

- When you close a file
- When the file handle's file system migrates
- When a client renames a file
- When the server reboots

If the client is unable to find the new file handle, an error message is logged in the `syslog` file. Further attempts to access this file fail with an I/O error.

Client Recovery in NFS Version 4

The NFS Version 4 protocol is a stateful protocol. Both the NFS client and the NFS server maintain current information about the open files and file locks.

When a server crashes and is rebooted, the server loses its state. The client detects that the server has rebooted and begins the process of helping the server re-establish the open and lock states that existed prior to the failure. This process is known as *client recovery* because the client directs the process.

When the client detects that the server has rebooted, the client immediately suspends its current activity and begins the process of client recovery. When the recovery process starts, a message such as the following is displayed in the system error log `/var/adm/messages`:

```
NOTICE: Starting recovery server server-name
```

During the recovery process, the client sends the server information about the client's previous state. However, during this period, the client does not send any new requests to the server. Any new requests to open files or set file locks must wait for the server to complete its recovery process before proceeding.

When the client recovery process is complete, the following message is displayed in the system error log `/var/adm/messages`:

```
NOTICE: Recovery done for server server-name
```

At this point, the client has successfully completed sending its state information to the server. However, even though the client has completed this process, other clients might not have done so. Therefore, for a period of time, known as the *grace period*, the server does not accept any open or lock requests to enable all clients to complete their recovery.

During the grace period, if the client attempts to open any new files or establish any new locks, the server denies the request with the `GRACE` error code. Upon receiving this error, the client must wait for the grace period to end and then resend the request to the server. During the grace period, the following message is displayed:

NFS server recovering

During the grace period, the commands that do not open files or set file locks can proceed. For example, the commands `ls` and `cd` do not open a file or set a file lock, these commands are not suspended. However, a command such as `cat`, which opens a file, would be suspended until the grace period ends.

When the grace period has ended, the following message is displayed:

```
NFS server recovery ok.
```

The client can now send new open and lock requests to the server.

Client recovery can fail for a variety of reasons. For example, if a network partition exists after the server reboots, the client might not be able to re-establish its state with the server before the grace period ends. When the grace period has ended, the server does not permit the client to re-establish its state because new state operations could create conflicts. For example, a new file lock might conflict with an old file lock that the client is trying to recover. When such situations occur, the server returns the `NO_GRACE` error code to the client.

If the recovery of an open operation for a file fails, the client marks the file as unusable and the following message is displayed:

```
WARNING: The following NFS file could not be recovered and was marked dead
(can't reopen: NFS status n): file : filename
```

If re-establishing a file lock during recovery fails, the following error message is displayed:

```
NOTICE: nfs4_send_siglost: pid process-ID lost
lock on server server-name
```

In this situation, the `SIGLOST` signal is posted to the process. The default action for the `SIGLOST` signal is to terminate the process.

To recover from this state, you must restart any applications that had files open at the time of the failure. Some processes that did not reopen the file could receive I/O errors. Other processes that did reopen the file, or performed the open operation after the recovery failure, can access the file without any problems.

Thus, some processes can access a particular file while other processes cannot.

OPEN Share Support in NFS Version 4

The NFS Version 4 protocol provides several file-sharing modes that the client can use to control file access by other clients. A client can specify the following:

- `DENY_NONE` mode permits other clients read and write access to a file.

- DENY_READ mode denies other clients read access to a file.
- DENY_WRITE mode denies other clients write access to a file.
- DENY_BOTH mode denies other clients read and write access to a file.

The Oracle Solaris NFS Version 4 server fully implements these file-sharing modes. Therefore, if a client attempts to open a file in a way that conflicts with the current share mode, the server denies the attempt by failing the operation. When such attempts fail with the initiation of the open or create operations, the NFS Version 4 client receives a protocol error. This error is mapped to the application error EACCES.

Even though the protocol provides several sharing modes, the open operation in Oracle Solaris does not offer multiple sharing modes. When opening a file, an Oracle Solaris NFS Version 4 client can only use the DENY_NONE mode.

Note - Even though the `fcntl` system call has an `F_SHARE` command to control file sharing, the `fcntl` commands cannot be implemented correctly with NFS Version 4. If you use these `fcntl` commands on an NFS Version 4 client, the client returns the `EAGAIN` error to the application.

Delegation in NFS Version 4

NFS Version 4 provides both client support and server support for delegation. *Delegation* is a technique by which the NFS server delegates the management of a file to a NFS client. For example, the server could grant either a read delegation or a write delegation to a client. Because read delegations do not conflict with each other, they can be granted to multiple clients at the same time. A write delegation can be granted to only one client because a write delegation conflicts with any file access by any other client. While holding a write delegation, the client does not send various operations to the server because the client is guaranteed exclusive access to a file. Similarly, the client does not send various operations to the server while holding a read delegation. Because the server guarantees that no client can open the file in write mode.

The effect of delegation is to greatly reduce the interactions between the server and the client for delegated files. Therefore, network traffic is reduced, and performance on the client and the server is improved. However, the degree of performance improvement depends on the kind of file interaction used by an application and the amount of network and server congestion.

A client does not request a delegation. The decision about whether to grant a delegation is made entirely by the server based on the access patterns for a file. If a file has been recently accessed in write mode by several different clients, the server might not grant a delegation because this access pattern indicates the potential for future conflicts.

A conflict occurs when a client accesses a file in a manner that is inconsistent with the delegations that are currently granted for that file. For example, if a client holds a write delegation on a file and a second client opens that file for read or write access, the server recalls the first client's write delegation. Similarly, if a client holds a read delegation and another client opens the same file for writing, the server recalls the read delegation. In both situations, the second client is not granted a delegation because a conflict now exists.

When a conflict occurs, the NFS server uses a callback mechanism to contact the NFS client that currently holds the delegation. Upon receiving this callback, the client sends the file's updated state to the server and returns the delegation. If the client fails to respond to the recall, the server revokes the delegation. In such instances, the server rejects all operations from the client for this file, and the client reports the requested operations as failures. Generally, these failures are reported to the application as I/O errors. To recover from these errors, the file must be closed and then reopened. Failures from revoked delegations can occur when a network partition exists between the client and the server while the client holds a delegation.

Note that one server cannot resolve access conflicts for a file that is stored on another server. Thus, an NFS server only resolves conflicts for files that it stores. Furthermore, in response to conflicts that are caused by clients that are running various versions of NFS, an NFS server can initiate recalls only to the client that is running NFS Version 4. An NFS server cannot initiate recalls for clients that are running earlier versions of NFS.

The process for detecting conflicts varies. For example, unlike NFS Version 4, because NFS Version 2 and NFS Version 3 do not have an open procedure, the conflict is detected only after the client attempts to read, write, or lock a file. The server's response to these conflicts varies also. For example:

- For NFS Version 3, the server returns the `JUKEBOX` error, which causes the client to halt the access request and try again later. The client displays the message `File unavailable`.
- For NFS Version 2, because an equivalent of the `JUKEBOX` error does not exist, the server makes no response, which causes the client to wait and then try again. The client displays the message `NFS server not responding`.

The error messages clear when the delegation conflict is resolved.

By default, NFS server delegation is enabled. You can disable delegation by setting the `server_delegation` parameter to `off`.

```
$ sharectl set -p server_delegation=off nfs
```

No keywords are required for NFS client delegation. The NFS Version 4 callback daemon, `nfs4cbd`, provides the callback service on the client. This daemon is started automatically whenever a mount for NFS Version 4 is enabled. By default, the client provides the necessary callback information to the server for all Internet transports that are listed in the `/etc/netconfig` system file. If the client is enabled for IPv6 and if the IPv6 address for the client's name can be determined, then the callback daemon accepts IPv6 connections.

The callback daemon uses a transient program number and a dynamically assigned port number. This information is provided to the server, and the server tests the callback path before granting any delegations. If the callback path does not test successfully, the server does not grant delegations, which is the only externally visible behavior.

Because callback information is embedded within an NFS Version 4 request, the server is unable to contact the client through a network device that uses Network Address Translation (NAT). Also, the callback daemon uses a dynamic port number. Therefore, the server might not be able to traverse a firewall, even if that firewall enables normal NFS traffic on port 2049. In such situations, the server does not grant delegations.

ACLs and `nfsmapid` in NFS Version 4

An access control list (ACL) provides file security by enabling the owner of a file to define file permissions for the file owner, the group, and other specific users and groups. On ZFS file systems, you can set ACLs on the NFS server and the NFS client by using the `chmod` command. For UFS file systems, you can use the `setfacl` command. For more information, see the [`chmod\(1\)`](#) and [`setfacl\(1\)`](#) man pages. In NFS Version 4, the ID mapper, `nfsmapid`, is used to map user IDs or group IDs in ACL entries on a server to user IDs or group IDs in ACL entries on a client. The reverse is also true: The user and group IDs in the ACL entries must exist on both the client and the server.

For more information about ACLs and `nfsmapid`, see the following:

- [“Setting ACLs on ZFS Files” in *Securing Files and Verifying File Integrity in Oracle Solaris 11.3*](#)
- [“NFS Daemons” on page 160](#)

ID Mapping Problems

The following situations can cause ID mapping to fail:

- If a user or group that exists in an ACL entry on the server cannot be mapped to a valid user or group on the client, the user can read the ACL but some of the users or groups will be shown as unknown.
For example, in this situation when you issue the `ls -lv` or `ls -lV` command, some of the ACL entries will have the group or user displayed as unknown.
- If the user ID or group ID in any ACL entry that is set on the client cannot be mapped to a valid user ID or group ID on the server, the `setfacl` and `chmod` commands can fail and return the `Permission denied` error message.

- If the client and server have mismatched `nfsmapid_domain` values, ID mapping fails. For more information, see [“NFS Daemons” on page 160](#).

To avoid ID mapping problems, do the following:

- Make sure that the value for `nfsmapid_domain` is set correctly. The currently selected NFSv4 domain is available in the `/var/run/nfs4_domain` file.
- Make sure that all user IDs and group IDs in the ACL entries exist on both the NFS Version 4 client and server.

Checking for Unmapped User IDs or Group IDs

To determine whether any user or group cannot be mapped on the NFS server or client, use the following script:

```
#!/usr/sbin/dtrace -Fs

sdt::nfs4-acl-nobody
{
    printf("validate_idmapping: (%s) in the ACL could not be mapped!",
stringof(arg0));
}
```

Note - The probe name that is used in this script is an interface that could change in the future. For more information, see [“Stability Levels” in *Oracle Solaris 11.3 DTrace \(Dynamic Tracing\) Guide*](#).

UDP and TCP Negotiation

In NFS Version 2 and NFS Version 3, negotiation for transport protocol happens at mount time. During initiation, the transport protocol is also negotiated. By default, the first connection-oriented transport that is supported on both the NFS client and the NFS server is selected. If this selection does not succeed, the first available connectionless transport protocol is used. The transport protocols that are supported on a system are listed in the `/etc/netconfig` file. TCP is the connection-oriented transport protocol that is supported by the release. UDP is the connectionless transport protocol.

When both the NFS protocol version and the transport protocol are determined by negotiation, the NFS protocol version is given precedence over the transport protocol. The NFS Version 3 protocol that uses UDP is given higher precedence than the NFS Version 2 protocol that is using

TCP. You can manually select both the NFS protocol version and the transport protocol with the `mount` command. For information about the NFS specific options for the `mount` command, see the [mount_nfs\(1M\)](#) man page. Under most conditions, allow the negotiation to select the best options.

File Transfer Size Negotiation

The file transfer size establishes the size of the buffers that are used when data is transferred between the NFS client and the NFS server. In general, larger transfer sizes are better. The NFS Version 3 protocol has an unlimited transfer size. Although the client can bid a smaller transfer size at mount time, under most conditions, this bid is not necessary.

The transfer size is not negotiated with systems that use the NFS Version 2 protocol. The maximum transfer size is set to 8 KB.

You can use the `-rsize` and `-wsize` options to manually set the transfer size with the `mount` command. You might need to reduce the transfer size for some system clients. Also, you can increase the transfer size if the NFS server is configured to use larger transfer sizes.

Note - Starting with the Oracle Solaris 10 release, restrictions on wire transfer sizes have been relaxed. The transfer size is based on the capabilities of the underlying transport. For example, the NFS transfer limit for UDP is still 32 KB. However, because TCP is a streaming protocol without the datagram limits of UDP, maximum transfer sizes over TCP have been increased to 1 MB.

How File Systems Are Mounted in NFS Version 3

The information in this section applies to NFS Version 3 mounts. The NFS Version 4 mount process does not include the portmap service or the MOUNT protocol.

When an NFS client attempts to mount a file system from an NFS server, the client must obtain a file handle from the server. The file handle must correspond to the file system. This process requires that several transactions occur between the client and the server. In this example, the NFS client is attempting to mount `/home/user` from the NFS server. A snoop trace for this transaction follows:

```
client -> server PORTMAP C GETPORT prog=100005 (MOUNT) vers=3 proto=UDP
server -> client PORTMAP R GETPORT port=33482
```

```
client -> server MOUNT3 C Null
server -> client MOUNT3 R Null
client -> server MOUNT3 C Mount /export/home9/user
server -> client MOUNT3 R Mount OK FH=9000 Auth=unix
client -> server PORTMAP C GETPORT prog=100003 (NFS) vers=3 proto=TCP
server -> client PORTMAP R GETPORT port=2049
client -> server NFS C NULL3
server -> client NFS R NULL3
client -> server NFS C FSINFO3 FH=9000
server -> client NFS R FSINFO3 OK
client -> server NFS C GETATTR3 FH=9000
server -> client NFS R GETATTR3 OK
```

In this trace, the NFS client first requests the mount port number from the portmap service on the NFS server. After the client receives the mount port number (33492), that number is used to test the availability of the service on the server. After the client has determined that a service is running on that port number, the client then makes a mount request. When the server responds to this request, the server includes the file handle for the file system (9000) being mounted. The client then sends a request for the NFS port number. When the client receives the number from the server, the client tests the availability of the NFS service (nfsd). Also, the client requests NFS information about the file system that uses the file handle.

In the following trace, the client is mounting the file system with the `public` option:

```
client -> server NFS C LOOKUP3 FH=0000 /export/home9/user
server -> client NFS R LOOKUP3 OK FH=9000
client -> server NFS C FSINFO3 FH=9000
server -> client NFS R FSINFO3 OK
client -> server NFS C GETATTR3 FH=9000
server -> client NFS R GETATTR3 OK
```

By using the default `public` file handle (which is `0000`), all the transactions to obtain information from the portmap service and to determine the NFS port number are skipped.

Note - NFS Version 4 provides support for volatile file handles. For more information, see [“Volatile File Handles in NFS Version 4” on page 32](#).

Effects of the `-public` Option and NFS URLs When Mounting

Using the `-public` option can create conditions that cause a mount to fail. Adding an NFS URL can also cause failures. The following list describes how a file system is mounted when you use these options:

- **Public option with NFS URL** – Use the public file handle. The mount fails if the public file handle is not supported.
- **Public option with regular path** – Use the public file handle. The mount fails if the public file handle is not supported.
- **NFS URL only** – Use the public file handle if this file handle is enabled on the NFS server. If the mount fails when you use the public file handle, then try the mount with the MOUNT protocol.
- **Regular path only** – Do not use the public file handle. The MOUNT protocol is used.

Client-Side Failover

Failover is process of selecting a server from a list of servers that support a replicated file system. Normally, the next server in the sorted list is used unless it fails to respond. By using client-side failover, an NFS client can detect when multiple servers are making the same data available and can switch to an alternative server when the current server is unavailable. This switch is known as remapping. Through normal use, the clients store the path name for each active file on a remote file system. During the remapping, these path names are evaluated to locate the files on the new server.

The file system can become unavailable if one of the following conditions occurs:

- If the file system is connected to a server that crashes
- If the server is overloaded
- If a network fault occurs

The failover under these conditions is normally transparent to the user. It can occur at any time without disrupting the processes that are running on the client.

For failover to occur, the file systems must be mounted read-only. The file systems must be identical for the failover to occur successfully. For information about identical file systems, see [“What Is a Replicated File System?” on page 42](#). A static file system or a file system that is not changed often is the best candidate for failover.

You cannot use CacheFS and client-side failover on the same NFS mount. Extra information is stored for each CacheFS file system. This information cannot be updated during failover, so you can use only one of these two features when mounting a file system.

The number of replicas that must be established for every file system depends on many factors. Ideally, you have a minimum of two servers. Each server supports multiple subnets. This setup is better than having a unique server on each subnet. The process requires the checking of each listed server. Therefore, if more servers are listed, each mount is slower.

What Is a Replicated File System?

For the purposes of client-side failover, a file system can be called a *replica* when it is the same size and has the same file size or file type as the original file system. Permissions, creation dates, and other file attributes are not considered. If the file size or file types are different, the remapping fails and the process hangs until the old server becomes available. In NFS Version 4, the behavior is different. For more information about client-side failover, see [“Client-Side Failover in NFS Version 4” on page 42](#).

You can maintain a replicated file system by using `rsync`, `cpio`, or another file transfer mechanism. Because updating the replicated file systems causes inconsistency, for best results consider these precautions:

- Rename the old version of the file before installing a new version of the file.
- Run the updates at night when client usage is low.
- Keep the updates small.
- Minimize the number of copies of the file.

Failover and NFS Locking

Some software packages require read locks on files. To prevent these products from breaking, read locks on read-only file systems are allowed but are visible to the client side only. The locks persist through a remapping because the server cannot detect the locks. Because the files should not change, you do not need to lock the file on the server side.

Client-Side Failover in NFS Version 4

In NFS Version 4, if a replica cannot be established because the file sizes are different or the file types are not the same, then the following happens:

1. The file is marked dead.
2. A warning is displayed.
3. The application that uses a file on the replicated mount receives a system call failure.

Note - If you restart the application and try again to access the file, you should be successful.

In NFS Version 4, you no longer receive replication errors for directories of different sizes. In prior versions of NFS, this condition was treated as an error and would impede the remapping process.

Furthermore, in NFS Version 4, if a directory read operation is unsuccessful, the operation is performed by the next listed server. In previous versions of NFS, unsuccessful read operations would cause the remapping to fail and the process to hang until the original server was available.

How NFS Server Logging Works

Note - Server logging is not supported in NFS Version 4.

NFS server logging provides records of NFS reads and writes, as well as operations that modify a file system. These records can be used to track access to information. In addition, the records can provide a quantitative way to measure interest in the information.

When a file system with logging enabled is accessed, the kernel writes raw data into a buffer file. This data includes the following:

- Time stamp
- Client IP address
- UID of the requester
- File handle of the file or directory object that is being accessed
- Type of operation that occurred

The `nfslogd` daemon converts this raw data into ASCII records that are stored in log files. During the conversion, the IP addresses are modified to host names and the UIDs are modified to logins if the name service that is enabled can find matches. The file handles are also converted into path names. To accomplish the conversion, the daemon tracks the file handles and stores information in a separate file handle-to-path table. That way, the path does not have to be identified again each time a file handle is accessed. Because no changes to the mappings are made in the file handle-to-path table if `nfslogd` is disabled, you must keep the daemon running.

How the WebNFS Service Works

The WebNFS service makes files in a directory available to clients by using a public file handle. A file handle is an address that is generated by the kernel that identifies a file for NFS clients. The *public file handle* has a predefined value, so the server does not need to generate a file handle for the client. The ability to use this predefined file handle reduces network traffic by eliminating the MOUNT protocol. This ability should also accelerate processes for the clients.

By default, the public file handle on an NFS server is established on the root file system. This default provides WebNFS access to any clients that already have mount privileges on the server. You can change the public file handle to point to any file system by using the share command.

When the NFS client has the file handle for the file system, a LOOKUP is run to determine the file handle for the file to be accessed. The NFS protocol allows the evaluation of only one path name component at a time. Each additional level of directory hierarchy requires another LOOKUP. A WebNFS server can evaluate an entire path name with a single multicomponent lookup transaction when the LOOKUP is relative to the public file handle. Multicomponent lookup enables the WebNFS server to deliver the file handle to the desired file without exchanging the file handles for each directory level in the path name.

In addition, an NFS client can initiate concurrent downloads over a single TCP connection. This connection provides quick access without the additional load on the server that is caused by setting up multiple connections. Although web browser applications support concurrent downloading of multiple files, each file has its own connection. By using one connection, the WebNFS software reduces the overhead on the server.

If the final component in the path name is a symbolic link to another file system, the client can access the file if the client already has access through normal NFS activities.

Normally, an NFS URL is evaluated relative to the public file handle. To change the evaluation to be relative to the server's root file system, add an additional slash to the beginning of the path. The following two NFS URLs are equivalent if the public file handle has been established on the /export/ftp file system.

```
nfs://server/junk
nfs://server//export/ftp/junk
```

Note - The NFS Version 4 protocol is preferred over the WebNFS service. NFS Version 4 fully integrates all the security negotiation that was added to the MOUNT protocol and the WebNFS service.

How WebNFS Security Negotiation Works

The NFS service includes a protocol that enables a WebNFS client to negotiate a selected security mechanism with a WebNFS server. The new protocol uses security negotiation multicomponent lookup, which is an extension to the multicomponent lookup that was used in earlier versions of the WebNFS protocol.

The WebNFS client initiates the process by making a regular multicomponent lookup request by using the public file handle. Because the client has no knowledge of how the path is protected

by the server, the default security mechanism is used. If the default security mechanism is not sufficient, the server replies with an AUTH_TOOWEAK error. The client needs to use a stronger default mechanism.

When the client receives the AUTH_TOOWEAK error, the client sends a request to the server to determine which security mechanisms are required. If the request succeeds, the server responds with an array of security mechanisms that are required for the specified path. Depending on the size of the array of security mechanisms, the client might have to make more requests to obtain the complete array. If the server does not support WebNFS security negotiation, the request fails.

After a successful request, the WebNFS client selects the first security mechanism from the array that the client supports. The client then issues a regular multicomponent lookup request by using the selected security mechanism to acquire the file handle. All subsequent NFS requests are made by using the selected security mechanism and the file handle.

WebNFS Limitations With Web Browser Use

The WebNFS software does not support several functions that a web site that uses HTTP can provide. These differences stem from the fact that the NFS server sends only the file, so any special processing must be done on the client. If you need to have one web site configured for both WebNFS and HTTP access, consider the following issues:

- NFS browsing does not run CGI scripts. So, a file system with an active web site that uses many CGI scripts might not be appropriate for NFS browsing.
- Accessing these files in different file formats through an NFS URL starts an external viewer if the file type can be determined by the file name. Because the WebNFS software does not check inside the file to determine the file type, the only way to determine a file type is by the file name extension. The browser should recognize any file name extension for a standard MIME type.
- NFS browsing cannot use server-side image maps but it can use client-side image maps because the URLs are defined with the location. No additional response is required from the document server.

Secure NFS Systems

The NFS environment is a powerful and convenient way to share file systems on a network of different computer architectures and operating systems. However, the same features that make sharing file systems through NFS operation convenient also pose some security problems. Historically, most NFS implementations have used UNIX (or AUTH_SYS) authentication,

but stronger authentication methods such as AUTH_DH have also been available. When using UNIX authentication, an NFS server authenticates a file request by authenticating the computer that makes the request but not the user. Therefore, an NFS client user can run `su` to become superuser and impersonate the owner of a file. If DH authentication is used, the NFS server authenticates the user, making this sort of impersonation much harder.

With root access and a knowledge of network programming, anyone can introduce arbitrary data into the network and extract any data from the network. The most dangerous attacks involve the introduction of data. An example is the impersonation of a user by generating the right packets or by recording "conversations" and replaying them later. These attacks affect data integrity. Attacks that involve passive eavesdropping, which is merely listening to network traffic without impersonating anybody, are not as dangerous because data integrity is not compromised. Users can protect the privacy of sensitive information by encrypting data that is sent over the network.

A common approach to network security problems is to leave the solution to each application. A better approach is to implement a standard authentication system at a level that covers all applications.

The Oracle Solaris operating system includes an authentication system at the level of the RPC, which is the mechanism on which the NFS operation is built. This system, known as Secure RPC, greatly improves the security of network environments and provides additional security to services such as the NFS system. An NFS system that uses the facilities that are provided by Secure RPC is known as a Secure NFS system.

Secure RPC

Secure RPC is fundamental to the Secure NFS system. The goal of Secure RPC is to build a system that is at minimum as secure as a time-sharing system. In a time-sharing system all users share a single computer and users are authenticated through login passwords. With Data Encryption Standard (DES) authentication, the same authentication process is completed. Users can log in on any remote computer just as users can log in on a local terminal. The users' login passwords are their assurance of network security. In a time-sharing environment, the system administrator has an ethical obligation not to change a password to impersonate someone. In Secure RPC, the network administrator is trusted not to alter entries in a database that stores *public keys*.

An RPC authentication system uses credentials and verifiers. Using ID badges as an example, the credential is what identifies a person: a name, address, and birthday. The verifier is the photo that is attached to the badge. You can be sure that the badge has not been stolen by checking the photo on the badge against the person who is carrying the badge. In RPC, the client process sends both a credential and a verifier to the server with each RPC request. The server sends back only a verifier because the client already "knows" the server's credentials.

RPC's authentication is open ended, which means that a variety of authentication systems can be plugged into it, such as UNIX, DH, and KERB.

When UNIX authentication is used by a network service, the credentials contain the client's host name, UID, GID, and group-access list. However, because no verifier exists, a superuser could falsify appropriate credentials by using commands such as `su`. Another problem is that UNIX authentication assumes all computers on a network are UNIX computers. UNIX authentication breaks down when applied to other operating systems in a heterogeneous network.

To overcome the problems of UNIX authentication, Secure RPC uses DH authentication.

Note - Although support for the Kerberos authentication system is no longer supplied as part of Secure RPC, a server-side and client-side implementation is included in the release. For more information about the implementation of Kerberos authentication, see [Chapter 2, “Kerberos on Oracle Solaris”](#) in *Managing Kerberos and Other Authentication Services in Oracle Solaris 11.3*.

DH Authentication

DH authentication uses the Data Encryption Standard (DES) and Diffie-Hellman public-key cryptography to authenticate both users and computers in the network. DES is a standard encryption mechanism. Diffie-Hellman public-key cryptography is a cipher system that involves two keys: one public and one secret. The public keys and secret keys are stored in the namespace. NIS stores the keys in the public-key map. These maps contain the public key and secret key for all potential users. For more information about how to set up the maps, see the [Working With Oracle Solaris 11.3 Directory and Naming Services: DNS and NIS](#).

The security of DH authentication is based on a sender's ability to encrypt the current time, which the receiver can then decrypt and check against its own clock. The timestamp is encrypted with DES. The two agents must agree on the current time, and sender and receiver must be using the same encryption key.

If a network runs a time-synchronization program, the time on the NFS client and the NFS server is synchronized automatically. If a time-synchronization program is not available, timestamps can be computed by using the server's time instead of the network time. The client asks the server for the time before starting the RPC session, then computes the time difference between its own clock and the server's. This difference is used to offset the client's clock when computing timestamps. If the client and server clocks become unsynchronized, the server begins to reject the client's requests. The DH authentication system on the client resynchronizes with the server.

The client and server arrive at the same encryption key by generating a random *conversation key*, also known as the *session key*, and by using public-key cryptography to deduce a *common*

key. The common key is a key that only the client and server are capable of deducing. The conversation key is used to encrypt and decrypt the client's timestamp. The common key is used to encrypt and decrypt the conversation key.

Using Secure RPC With NFS

Be aware of the following points if you plan to use Secure RPC:

- If an NFS server crashes when no system administrator is available (after a power failure, for example), all the secret keys that are stored on the system are deleted. No process can access secure network services or mount an NFS file system. The important processes during a reboot are usually run as `root`. Therefore, these processes would work if `root`'s secret key were stored away, but nobody is available to type the password that decrypts it. `keylogin -r` allows `root` to store the clear secret key in `/etc/.rootkey`, which `keyserv` reads.
- Some systems boot in single-user mode, with a `root` login shell on the console and no password prompt. Physical security is imperative in such cases.
- Diskless computer booting is not totally secure. Somebody could impersonate the boot server and boot a devious kernel that, for example, makes a record of the secret key on a remote computer. The Secure NFS system provides protection only after the kernel and the key server are running. Otherwise, no way exists to authenticate the replies that are given by the boot server. This limitation could be a serious problem, but the limitation requires a sophisticated attack using kernel source code. Also, the crime would leave evidence. If you polled the network for boot servers, you would discover the devious boot server's location.
- Most `setuid` programs are owned by `root`. If the secret key for `root` is stored in `/etc/.rootkey`, these programs behave as they always have. If a `setuid` program is owned by a user, however, the `setuid` program might not always work. For example, suppose that a `setuid` program is owned by `dave` and `dave` has not logged into the computer since it booted. The program would not be able to access secure network services.
- If you log in to a remote computer (using `login`, `rlogin`, or `telnet`) and use `keylogin` to gain access, you provide access to your account. Your secret key is passed to that computer's key server, which then stores your secret key. This process is only a concern if you do not trust the remote computer. If you have doubts, however, do not log in to a remote computer if the remote computer requires a password. Instead, use the NFS environment to mount file systems that are shared by the remote computer. As an alternative, you can use `keylogout` to delete the secret key from the key server.
- If a home directory is shared with the `-o sec=dh` option, remote logins can be a problem. If the `/etc/hosts.equiv` or `~/.rhosts` files are not set to prompt for a password, the login succeeds. However, the users cannot access their home directories because no authentication has occurred locally. If users are prompted for a password, they have access to their home directories if the password matches the network password.

How Mirror Mounts Work

The Oracle Solaris 11 release includes a new mounting facility called *mirror mounts*. Mirror mounts allow an NFS Version 4 client to access files in a file system as soon as the file system is shared on an NFS Version 4 server. The files can be accessed without the overhead of using the `mount` command or updating `autofs` maps. In effect, after an NFS Version 4 file system is mounted on a client, any other file systems from that server can also be mounted.

Generally, using the mirror mounts facility is optimal for your NFS Version 4 clients except when you need to do the following:

- Use a different hierarchy on the client than exists on the server
- Use different mount options than those of the parent file system

Mounting a File System Using Mirror Mounts

If a file system is mounted on an NFS Version 4 client by using manual mounts or `autofs`, any additional file systems that are added to the mounted file system can be mounted on the client by using the mirror mount facility. The client requests access to the new file system with the same mount options that were used on the parent directory. If the mount fails for any reason, the normal NFS Version 4 security negotiations occur between the server and the client to adjust the mount options so that the mount request succeeds.

When an automount trigger exists for a particular server file system, the automount trigger takes precedence over mirror mounting, so a mirror mount will not occur for that file system. To use mirror mounts in this case, the automount entry must be removed.

In the Oracle Solaris 11 release, accessing the `/net` or `/home` automount point causes a mount of the `/net` or `/home` server namespace. Access to directories or files under those directories is given through the mirror mounts facility.

For specific instructions about how to use mirror mounts, see [“How to Mount All File Systems From a Server” on page 77](#).

Unmounting a File System Using Mirror Mounts

Mirror-mounted file systems are automatically unmounted if they are idle after a certain period of inactivity. The period is set by using the `timeout` parameter, which is used by the automounter for the same purpose.

If an NFS file system is manually unmounted, then any mirror-mounted file systems contained within it are also unmounted, if idle. If a mirror-mounted file system is active, the manual unmount fails as though that original file system were busy. However, a forced unmount is propagated through all the contained mirror-mounted file systems.

If a file system boundary is encountered within an automounted file system, a mirror mount occurs. When the automounter unmounts the parent file system, any mirror-mounted file systems within it are also automatically unmounted, if idle. If there is an active mirror-mounted file system, the automatic unmount does not occur, which preserves current automount behavior.

How NFS Referrals Work

Starting with the Oracle Solaris 11.1 release, NFS referrals enable an NFS Version 4 server to point to file systems located on other NFS Version 4 servers as a way of connecting multiple NFS Version 4 servers into a uniform namespace.

NFS Version 2, NFS Version 3, and other types of clients can follow a referral because it appears to them to be a symbolic link.

When to Use NFS Referrals

NFS referrals are useful when you want to create what appears to be a single set of file names across multiple servers, and you prefer not to use `autofs` for this purpose. Note that only NFS Version 4 servers can be used and that the servers must be running at least the Oracle Solaris 11.1 release to host a referral.

Creating an NFS Referral

You create an NFS referral by using the `nfsref` command. When a referral is created and the mount point does not yet exist, a symbolic link is generated. This symbolic link includes a special flag that identifies an object as a reparse point. A reparse point is a special marker used to note that special handling is required. If the reparse point already exists, NFS service data is added or replaces existing NFS service data, as appropriate.

Removing an NFS Referral

You remove NFS referral by using the `nfsref` command. The command removes NFS service data from the specified reparse point. It also removes the reparse point if no other types of service data are present.

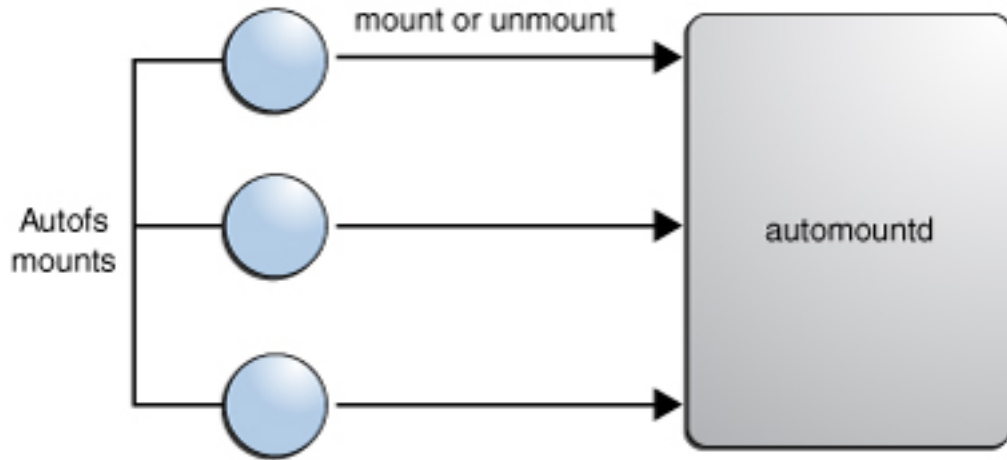
How Autofs Works

Autofs is a kernel file system that supports automatic mounting and unmounting. The components that work together to accomplish automatic mounting are:

- automount command
- autofs file system
- automountd daemon

Autofs is a client-side service that automatically mounts the appropriate file system. The automount service, `svc:/system/filesystem/autofs`, which is called at system startup time, reads the master map file `auto_master` to create the initial set of autofs mounts. These autofs mounts are not automatically mounted at startup time but rather are points under which file systems are mounted in the future. These points are also known as trigger nodes. For more information about starting the navigation process, see [“How Autofs Starts the Navigation Process” on page 59](#).

The following figure shows how the autofs service starts the automountd daemon.

FIGURE 3 `svc:/system/filesystem/autofs` Service Starts automount

After the autofs mounts are set up, these mounts can trigger file systems to be mounted under them. For example, when autofs receives a request to access a file system that is not currently mounted, autofs calls automountd, which actually mounts the requested file system.

When a request is made to access a file system at an autofs mount point, the following occurs:

1. Autofs intercepts the request.
2. Autofs sends a message to the automountd daemon for the requested file system to be mounted.
3. The automountd daemon locates the file system information in a map, creates the trigger nodes, and performs the mount.
4. Autofs allows the intercepted request to proceed.
5. Autofs unmounts the file system after a period of inactivity.

After initially mounting autofs mounts, use the `automount` command to update autofs mounts as necessary. The command compares the list of mounts in the `auto_master` map with the list of mounted file systems in the mount table file `/etc/mnttab` (formerly `/etc/mntab`). `automount` then makes the appropriate changes. This process enables system administrators to change mount information within `auto_master` and have those changes used by the autofs processes without stopping and restarting the autofs daemon. After the file system is mounted, further access does not require any action from automountd until the file system is automatically unmounted.

Unlike `mount`, `automount` does not read the `/etc/vfstab` file (which is specific to each computer) for a list of file systems to mount. The `automount` command is controlled within a domain and on computers through the namespace or local files.

Note - Mounts that are managed through the `autofs` service should not be manually mounted or unmounted. Even if the operation is successful, the `autofs` service does not check that the object has been unmounted, resulting in possible inconsistencies. A reboot clears all the `autofs` mount points.

How Autofs Navigates Through the Network

Autofs searches a series of maps to navigate through the network. Maps are files that contain information such as the password entries of all users on a network or the names of all host computers on a network. Effectively, the maps contain network-wide equivalents of UNIX administration files. Maps are available locally or through a network name service such as NIS and LDAP.

Autofs Maps

Autofs uses three types of maps:

- Master map
- Direct maps
- Indirect maps

Master Autofs Map

The `auto_master` map associates a directory with a map. The map is a master list that specifies all the maps that `autofs` should check. The following example shows the types of information that an `auto_master` file could contain.

EXAMPLE 1 Sample `/etc/auto_master` File

```
# Master map for automounter
#
+auto_master
/net          -hosts          -nosuid,nobrowse
```

```

/home          auto_home      -nobrowse
/nfs4         -fedfs         -ro,nosuid,nobrowse
/-           auto_direct    -ro

```

This example shows the generic `auto_master` file with one addition for the `auto_direct` map. Each line in the master map `/etc/auto_master` has the following syntax:

mount-point map-name [mount-options]

mount-point Full (absolute) path name of a directory. If the directory does not exist, `autofs` creates the directory if possible. If the directory exists and is not empty, mounting on the directory hides its contents. In this situation, `autofs` issues a warning.

The notation `/-` as a mount point indicates that this particular map is a direct map. The notation also means that no particular mount point is associated with the map.

map-name Name of the map that `autofs` uses to find directions to locations, or mount information. If the name is preceded by a slash (`/`), `autofs` interprets the name as a local file. Otherwise, `autofs` searches for the mount information by using the search that is specified in the name-service switch configuration file (`/etc/nsswitch.conf`). Special maps are also used for `/net`. For more information, see [“Mount Point /net” on page 55](#).

mount-options An optional, comma-separated list of options that apply to the mounting of the entries that are specified in *map-name*, unless the entries in *map-name* list other options. Options for each specific type of file system are listed in the `mount` man page for that file system. For information about NFS-specific mount options, see the `mount_nfs(1M)` man page. For NFS-specific mount points, the `bg` (background) and `fg` (foreground) options do not apply.

A line that begins with `#` is a comment. All the text that follows until the end of the line is ignored.

To split long lines into shorter ones, put a backslash (`\`) at the end of the line. The maximum number of characters of an entry is 1024.

Note - If the same mount point is used in two entries, the first entry is used by the automount command. The second entry is ignored.

Mount Point /home

The mount point `/home` is the directory under which the entries that are listed in `/etc/auto_home` (an indirect map) are to be mounted.

Note - Autofs runs on all computers and supports `/net` and `/home` (automounted home directories) by default. You can override these default entries in the NIS `auto.master` map or by local editing of the `/etc/auto_master` file.

Mount Point /net

Autofs mounts under the directory `/net` all the entries in the special built-in map `-hosts` that uses only the hosts database. Suppose that the computer `system1` is in the hosts database and it exports any of its file systems. The following command changes the current directory to the root directory of the computer `gumbo`.

```
$ cd /net/gumbo
```

Autofs can mount only the *exported* file systems of host `system1`, that is, those file systems on a server that are available to network users instead of those file systems on a local disk. Therefore, all the files and directories on `system1` might not be available through `/net/system1`.

With the `/net` method of access, the server name is in the path and is location dependent. If you want to move an exported file system from one server to another, the path might no longer work. Instead, you should set up an entry in a map specifically for the file system you want rather than using `/net`.

Note - Using NFS Version 3 and earlier protocols, `autofs` checks the server's export list only at mount time. After a server's file systems are mounted, `autofs` does not check with the server again until the server's file systems are automatically unmounted. Therefore, newly exported file systems are not "seen" until the file systems on the client are unmounted and then remounted. For systems using NFS Version 4, mirror mounts reflect any dynamic changes made to the list of exported file systems on the server.

Mount Point /nfs4

The `/nfs4` mount point uses a pseudo-map to mount the FedFS domain root. A reference to the `/nfs4/example.net` file results in an attempt to find the domain root for the DNS

domain `example.net` and mounts it at that location. Mounting a path under `/nfs4` requires that the DNS server returns a record, as described in [“Set Up a DNS Record for a FedFS Server” on page 91](#).

Direct Autofs Maps

A direct map is an automount point. With a direct map, a direct association exists between a mount point on the client and a directory on the server. Direct maps have a full path name and indicate the relationship explicitly. The following example shows a typical `/etc/auto_direct` map:

```
/usr/local      -ro \  
  /bin          system1:/export/local/sun4 \  
  /share       system1:/export/local/share \  
  /src         system1:/export/local/src \  
/usr/man       -ro  system2:/usr/man \  
              system3:/usr/man \  
              system4:/usr/man \  
/usr/games     -ro  system5:/usr/games \  
/usr/spool/news -ro  system6:/usr/spool/news \  
              system4:/var/spool/news
```

Lines in direct maps have the following syntax:

key [*mount-options*] *location*

key Path name of the mount point in a direct map.

mount-options Options that you want to apply to this particular mount. These options are required only if the options differ from the map default. Options for each specific type of file system are listed in the `mount` man page for that file system. For information about NFS specific mount options, see the [`mount_nfs\(1M\)`](#) man page.

location Location of the file system. One or more file systems are specified as *server:pathname* for NFS file systems.

Note - The path name should not include an automounted mount point. The path name should be the actual absolute path to the file system. For instance, the location of a home directory should be listed as *server:/export/home/username*, not as *server:/home/username*.

As in the master map, a line that begins with `#` is a comment. All the text that follows until the end of the line is ignored. Put a backslash at the end of the line to split long lines into shorter ones.

Of all the maps, the entries in a direct map most closely resemble the corresponding entries in `/etc/vfstab`. An entry might appear in `/etc/vfstab` as follows:

```
dancer:/usr/local - /usr/local/tmp nfs - yes ro
```

The equivalent entry appears in a direct map as follows:

```
/usr/local/tmp -ro dancer:/usr/local
```

Note - No concatenation of options occurs between the automounter maps. Any options that are added to an automounter map override all options that are listed in maps that are searched earlier. For instance, options that are included in the `auto_master` map would be overridden by corresponding entries in any other map.

For information about the features of direct autofs map, see [“How Autofs Selects the Nearest Read-Only Files for Clients” on page 61](#).

Mount Point /-

In [Example 1, “Sample `/etc/auto_master` File,” on page 53](#), the mount point `/-` tells autofs not to associate the entries in `auto_direct` with any specific mount point. Indirect maps use mount points that are defined in the `auto_master` file. Direct maps use mount points that are specified in the named map. Note that, in a direct map the key, or mount point, is a full path name.

An NIS `auto_master` file can have only one direct map entry because the mount point must be a unique value in the namespace. An `auto_master` file that is a local file can have any number of direct map entries if entries are not duplicated.

Indirect Autofs Maps

An indirect map uses a substitution value of a key to establish the association between a mount point on the client and a directory on the server. Indirect maps are useful for accessing specific file systems, such as home directories. The `auto_home` map is an example of an indirect map.

Lines in indirect maps have the following general syntax:

```
key [mount-options] location
```

key Name without slashes in an indirect map.

| | |
|----------------------|---|
| <i>mount-options</i> | Options that you want to apply to this particular mount. These options are required only if the options differ from the map default. Options for each specific type of file system are listed in the mount man page for that file system. For example, see the mount_nfs(1M) man page for NFS-specific mount options. |
| <i>location</i> | Location of the file system. One or more file systems are specified as <i>server:pathname</i> . |

Note - The path name should not include an automounted mount point. The path name should be the actual absolute path to the file system. For instance, the location of a directory should be listed as *server:/usr/local*, not as *server:/net/server/usr/local*.

As in the master map, a line that begins with # is a comment. All the text that follows until the end of the line is ignored. Put a backslash (\) at the end of the line to split long lines into shorter ones. [Example 1, "Sample /etc/auto_master File," on page 53](#) shows an auto_master map that contains the following entry:

```
/home      auto_home      -nobrowse
```

auto_home is the name of the indirect map that contains the entries to be mounted under /home. A typical auto_home map might contain the following:

```
user1      server1:/export/home/user1
user2      server2:/export/home/user2
user3      server3:/export/home/user3
user4      server4:/export/home/user4
user5      server5:/export/home/user5
user6      server6:/export/home/user6
user7      -rw,nosuid  server7:/export/home/user7
```

As an example, assume that the previous map is on host master-server. Suppose that the user user7 has an entry in the password database that specifies her home directory as /home/user7. Whenever user7 logs in to computer master-server, autofs mounts the directory /export/home/user7 that resides on the computer server7. Her home directory is mounted read-write, nosuid.

Assume the following conditions occur: User user7's home directory is listed in the password database as /home/user7. Anybody, including user7, has access to this path from any computer that is set up with the master map referring to the auto_home map.

Under these conditions, user user7 can run login or rlogin on any of these computers and have her home directory mounted in place for her.

Furthermore, now user7 can also type the following command:

```
$ cd ~user1
```

Autofs mounts user1's home directory for user7 (if all permissions allow).

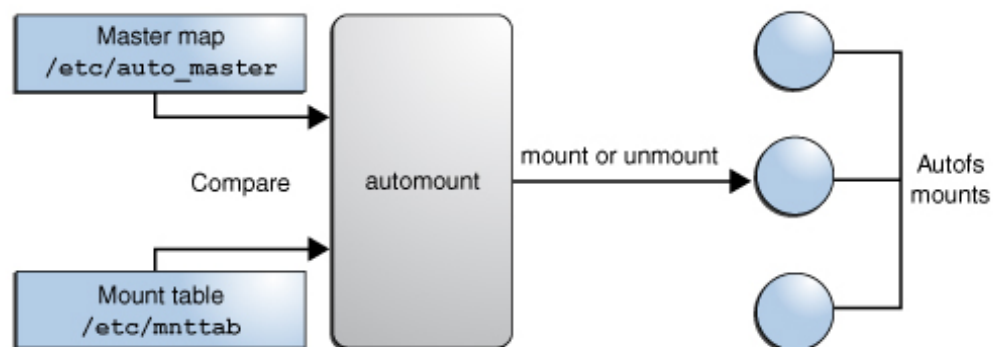
Note - No concatenation of options occurs between the automounter maps. Any options that are added to an automounter map override all options that are listed in maps that are searched earlier. For instance, options that are included in the `auto_master` map are overridden by corresponding entries in any other map.

On a network without a name service, you have to change all the relevant files (such as `/etc/passwd`) on all systems on the network to allow Linda access to her files. With NIS, make the changes on the NIS master server and propagate the relevant databases to the slave servers.

How Autofs Starts the Navigation Process

The `automount` command reads the master map at system startup. Each entry in the master map is a direct map name or an indirect map name, its path, and its mount options. The specific order of the entries is not important.

FIGURE 4 Navigation Through the Master Map



This figure shows that `automount` compares entries in the master map with entries in the mount table to generate a current list.

Autofs Mount Process

What the autofs service does when a mount request is triggered depends on how the automounter maps are configured. The mount process is generally the same for all mounts. However, the final result changes with the mount point that is specified and the complexity of the maps. The mount process includes the creation of the trigger nodes.

Simple Autofs Mount

To help explain the autofs mount process, assume that the following files are installed.

```
$ cat /etc/auto_master
# Master map for automounter
#
+auto_master
/net      -hosts      -nosuid,nobrowse
/home     auto_home   -nobrowse
/share    auto_share
$ cat /etc/auto_share
# share directory map for automounter
#
ws        gumbo:/export/share/ws
```

When the /share directory is accessed, the autofs service creates a trigger node for /share/ws, which is an entry in /etc/mnttab that resembles the following entry:

```
-hosts /share/ws  autofs  nosuid,nobrowse,ignore,nest,dev=###
```

When the /share/ws directory is accessed, the autofs service completes the process as follows:

1. Checks the availability of the server's mount service.
2. Mounts the requested file system under /share. Now the /etc/mnttab file contains the following entries.

```
-hosts /share/ws  autofs  nosuid,nobrowse,ignore,nest,dev=###
gumbo:/export/share/ws /share/ws  nfs  nosuid,dev=####  #####
```

Hierarchical Mounting

When multiple layers are defined in the automounter files, the mount process becomes more complex. Suppose that you expand the /etc/auto_shared file from the previous example to contain the following:

```
# share directory map for automounter
#
ws      /      gumbo:/export/share/ws
        /usr   gumbo:/export/share/ws/usr
```

The mount process is basically the same as the previous example when the `/share/ws` mount point is accessed. In addition, a trigger node to the next level (`/usr`) is created in the `/share/ws` file system so that the next level can be mounted if it is accessed. In this example, `/export/share/ws/usr` must exist on the NFS server for the trigger node to be created.



Caution - Do not use the `-soft` option when specifying hierarchical layers. For more information, see [“Autofs Unmounting” on page 61](#).

Autofs Unmounting

The unmounting that occurs after a certain amount of idle time is from the bottom up (reverse order of mounting). If one of the directories at a higher level in the hierarchy is busy, only file systems below that directory are unmounted. During the unmounting process, any trigger nodes are removed and then the file system is unmounted. If the file system is busy, the unmount fails and the trigger nodes are reinstalled.



Caution - Do not use the `-soft` option when specifying hierarchical layers. If the `-soft` option is used, requests to reinstall the trigger nodes can time out. The failure to reinstall the trigger nodes leaves no access to the next level of mounts. The only way to clear this problem is to have the automounter unmount all of the components in the hierarchy. The automounter can complete the unmounting either by waiting for the file systems to be automatically unmounted or by rebooting the system.

How Autofs Selects the Nearest Read-Only Files for Clients

This section uses the following example direct map to help explain how autofs selects the nearest read-only files for clients.

```
/usr/local      -ro \
  /bin          ivy:/export/local/sun4\
  /share        ivy:/export/local/share\
  /src          ivy:/export/local/src
/usr/man        -ro oak:/usr/man \
               rose:/usr/man \
```

```
willow:/usr/man
/usr/games      -ro peach:/usr/games
/usr/spool/news -ro pine:/usr/spool/news \
willow:/var/spool/news
```

The mount points `/usr/man` and `/usr/spool/news` list more than one location, with three locations for the first mount point and two locations for the second mount point. Any of the replicated locations can provide the same service to any user. This procedure is sensible only when you mount a file system that is read-only, as you must have some control over the locations of files that you write or modify. You want to avoid modifying files on one server on one occasion and, minutes later, modifying the "same" file on another server. The benefit is that the best available server is used automatically without any effort by the user.

If the file systems are configured as replicas (see [“What Is a Replicated File System?” on page 42](#)), the NFS clients have the advantage of using failover. Not only is the best NFS server automatically determined, but if that server becomes unavailable, the client automatically uses the next-best server.

An example of a good file system to configure as a replica is man pages. In a large network, more than one server can export the current set of man pages. Which server you mount the man pages from does not matter as long as the server is running and exporting its file systems. In the direct map example, multiple mount locations are expressed as a list of mount locations in the map entry.

```
/usr/man -ro oak:/usr/man rose:/usr/man willow:/usr/man
```

In this example, you can mount the man pages from the servers `oak`, `rose`, or `willow`. Which server is best depends on a number of factors, including the following:

- Number of servers that support a particular NFS protocol level
- Proximity of the server
- Weighting

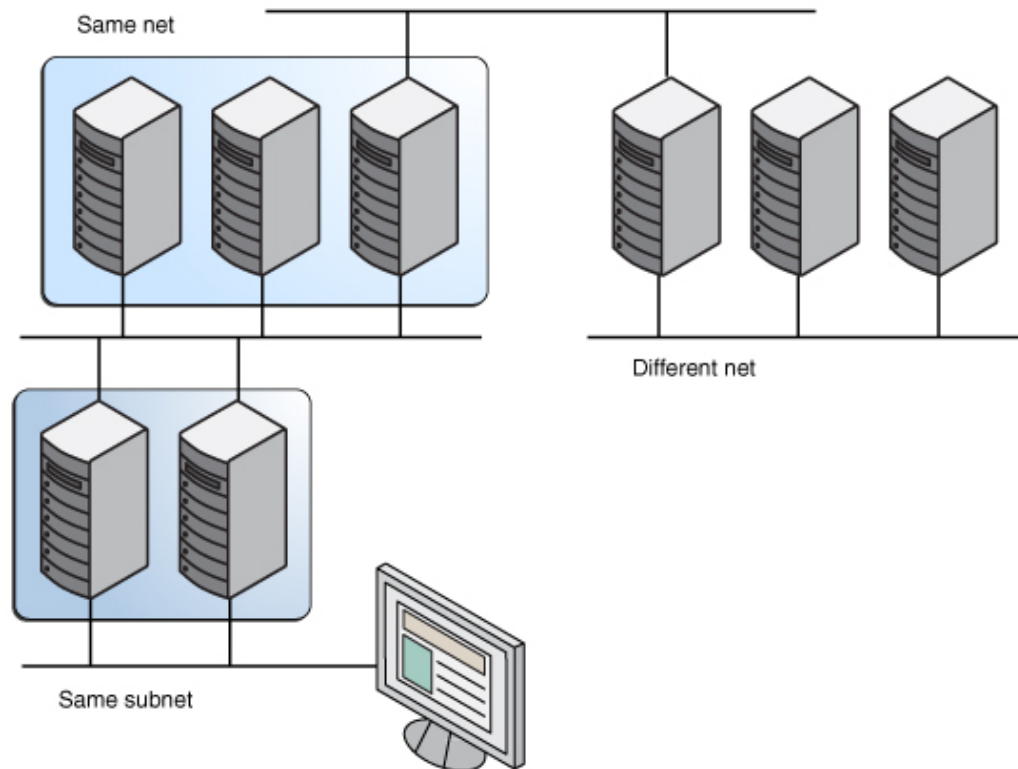
During the sorting process, a count is taken of the number of servers that support each version of the NFS protocol. Whichever version of the protocol is supported on the most servers becomes the protocol that is used by default. This selection provides the client with the maximum number of servers to depend on.

After the largest subset of servers with the same version of the protocol is found, that server list is sorted by proximity. To determine proximity, IPv4 addresses are inspected to determine which servers are in each subnet. Servers on a local subnet are given preference over servers on a remote subnet. Preference for the closest server reduces latency and network traffic.

Note - Proximity cannot be determined for replicas that are using IPv6 addresses.

Figure 5, [“Server Proximity,” on page 63](#) illustrates server proximity.

FIGURE 5 Server Proximity



If several servers that support the same protocol are on the local subnet, the time to connect to each server is determined and the fastest server is used. The sorting can also be influenced by using weighting. For more information about weighting, see [“Autofs and Weighting”](#) on page 65.

For example, if NFS Version 4 servers are more abundant on the local subnet, NFS Version 4 becomes the protocol that is used by default. However, the sorting process is more complex when servers on the local subnet support different protocols. Here are some examples of how the sorting process works:

- Servers on the local subnet are given preference over servers on a remote subnet. So, if a NFS Version 3 server is on the local subnet and the closest NFS Version 4 server is on a remote subnet, the NFS Version 3 server is given preference. Likewise, if the local subnet

consists of NFS Version 2 servers, they are given preference over remote subnets with NFS Version 3 and NFS Version 4 servers.

- If the local subnet consists of a varied number of NFS Version 2, NFS Version 3, and NFS Version 4 servers, more sorting is required. The automounter prefers the highest version on the local subnet. In this instance, NFS Version 4 is the highest version. However, if the local subnet has more NFS Version 3 or NFS Version 2 servers than NFS Version 4 servers, the automounter "bids down" from the highest version on the local subnet by one version. For example, if the local subnet has three servers with NFS Version 4, three servers with NFS Version 3, and ten servers with NFS Version 2, a NFS Version 3 server is selected.
- Similarly, if the local subnet consists of a varied number of NFS Version 2 and NFS Version 3 servers, the automounter first looks at the which version represents the highest version on the local subnet. Next, the automounter counts the number of servers that run each version. If the highest version on the local subnet also represents the most servers, the highest version is selected. If a lower version has more servers, the automounter bids down from the highest version on the local subnet by one version. For example, if more NFS Version 2 servers are on the local subnet than NFS Version 3 servers, a NFS Version 2 server is selected.

Note - Weighting is also influenced by parameters stored in the SMF repository. Specifically the values for `server_versmin`, `client_versmin`, `server_versmax` and `client_versmax` can exclude some versions from the sorting process. For more information about these parameters, see [“NFS Daemons” on page 160](#).

With failover, the sorting is checked at mount time when a server is selected. Multiple locations are useful in an environment where individual servers might not export their file systems temporarily.

Failover is particularly useful in a large network with many subnets. Autofs chooses the appropriate server and is able to confine NFS network traffic to a segment of the local network. If a server has multiple network interfaces, you can list the host name that is associated with each network interface as if the interface were a separate server. Autofs selects the nearest interface to the client.

Note - No weighting and no proximity checks are performed with manual mounts. The `mount` command prioritizes the servers that are listed from left to right.

For more information, see [`automount\(1M\)`](#) man page.

Autofs and Weighting

You can influence the selection of servers at the same proximity level by adding a weighting value to the autofs map. For example:

```
/usr/man -ro oak,rose(1),willow(2):/usr/man
```

The numbers in parentheses indicate a weighting. Servers without a weighting have a value of zero and, therefore, are most likely to be selected. The higher the weighting value, the lower the chance that the server is selected.

Note - All other server selection factors are more important than weighting. Weighting is only considered when selecting between servers with the same network proximity.

Variables in an Autofs Map Entry

You can create a client-specific variable by prefixing a dollar sign (\$) to its name. The variable helps you to accommodate different architecture types that are accessing the same file system location. You can also use curly braces to delimit the name of the variable from appended letters or digits. The following table shows the predefined map variables.

TABLE 1 Predefined Map Variables

| Variable | Meaning | Derived From | Example |
|----------|---|--------------|---------|
| ARCH | Architecture type | uname -m | sun4 |
| CPU | Processor type | uname -p | sparc |
| HOST | Host name | uname -n | system1 |
| OSNAME | Operating system name | uname -s | SunOS |
| OSREL | Operating system release | uname -r | 5.10 |
| OSVERS | Operating system version (version of the release) | uname -v | GENERIC |

You can use variables anywhere in an entry line except as a key. For instance, suppose that you have a file server that exports binaries for SPARC and x86 architectures from `/usr/local/bin/sparc` and `/usr/local/bin/x86` respectively. The clients can mount through a map entry such as the following:

```
/usr/local/bin -ro server:/usr/local/bin/$CPU
```

The same entry for all clients now applies to all architectures.

Note - Most applications that are written for any of the sun4 architectures can run on all sun4 platforms. The `-ARCH` variable is hard-coded to sun4.

Maps That Refer to Other Maps

Special characters used with map names in map entries in a file map affect how the map name is processed.

- A map entry `+mapname` that is used in a file map causes automount to read the specified map as if it were included in the current file.
- If `mapname` is not preceded by a slash, autofs treats the map name as a string of characters and uses the name-service switch policy to find the map name. If the path name is an absolute path name, automount checks a local map of that name.
- If the map name starts with a dash (`-`), automount consults the appropriate built-in map, such as `hosts`.

The `svc:system/name-service/switch` service contains the search order for the naming services. The `automount` property in the `config` property group specifies the order that the name service databases are searched when looking for automount entries. If no specific `config/automount` property is specified, then the order defined in the `config/default` property is used.

EXAMPLE 2 Displaying the Search Order of Maps by the `automount` command

```
$ svcprop -p config svc:/system/name-service/switch
config/value_authorization astring solaris.smf.value.name-service.switch
config/printer astring user\ files
config/default astring files\ nis
config/automount astring files\ nis
```

The example shows that the maps in the local files are searched before the NIS maps. The same would be true if the `config/automount` property was not specified because the `config/default` entry would be used. Therefore, you can have a few entries in your local `/etc/auto_home` map for the most commonly accessed home directories. You can then use the switch to fall back to the NIS map for other entries.

```
bill          cs.csc.edu:/export/home/bill
bonny         cs.csc.edu:/export/home/bonny
```

After consulting the included map, if no match is found, automount continues scanning the current map. Therefore, you can add more entries after a `+` entry.

```
bill          cs.csc.edu:/export/home/bill
bonny         cs.csc.edu:/export/home/bonny
```

```
+auto_home
```

The map that is included can be a local file or a built-in map. Only local files can contain + entries.

```
+/etc/auto_mystuff    # local map
+auto_home           # NIS map
+-hosts              # built-in hosts map
```

Note - You cannot use + entries in NIS maps.

Executable Autofs Maps

You can create an autofs map that executes some commands to generate the autofs mount points. Executable autofs maps are useful if you need to be able to create the autofs structure from a database or a flat file. The disadvantage to using an executable map is that the map needs to be installed on each host. An executable map cannot be included in NIS and LDAP name services.

The executable map must have an entry in the auto_master file.

```
/execute    auto_execute
```

The following example shows a sample executable map:

```
#!/bin/ksh
#
# executable map for autofs
#

case $1 in
    src) echo '-nosuid,hard bee:/export1' ;;
esac
```

For this example to work, the file must be installed as /etc/auto_execute and must have the executable bit set. Set permissions to 744. Under these circumstances, running the following command causes the /export1 file system from bee to be mounted:

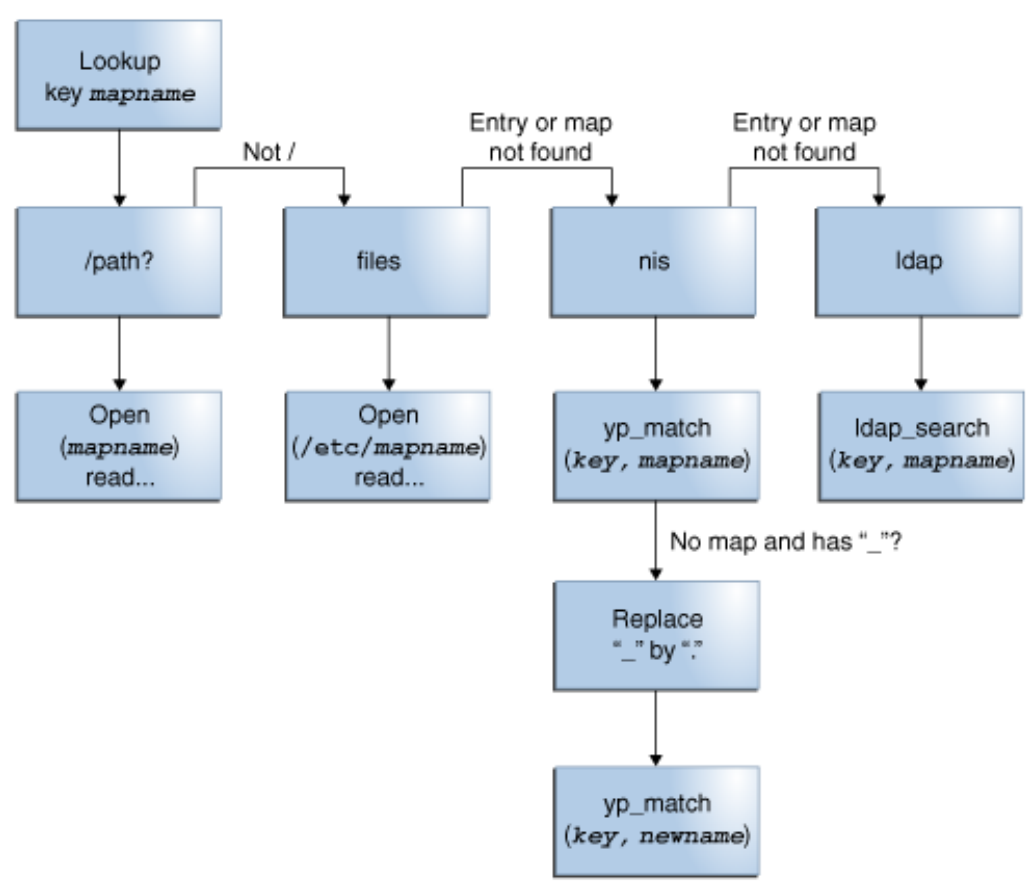
```
$ ls /execute/src
```

Default Autofs Behavior With Name Services

At boot time, autofs is invoked by the service svc:/system/filesystem/autofs and autofs checks for the master auto_master map.

Autofs uses the name service ordering specified in the `config/automount` property of the `svc:/system/name-service/switch` service. If the `config/automount` property is not defined, then the `config/default` property is used. If NIS is selected and autofs cannot find a map that autofs can use, but does find a map name that contains one or more underscores, the underscores are changed to dots to allow traditional NIS file names to work. Then autofs checks the map again. For a name service switch setting of `automount: files nis ldap`, maps would be found as shown in the following figure.

FIGURE 6 How Autofs Uses the Name Service



The screen activity for this session would resemble the following example.

```
$ grep /home /etc/auto_master
/home          auto_home
```

```
$ ypmatch brent auto_home
Can't match key brent in map auto_home. Reason: no such map in
server's domain.
```

```
$ ypmatch brent auto.home
diskus:/export/home/diskus1/&
```

If "files" is selected as the name service, all maps are assumed to be local files in the /etc directory. Autofs interprets a map name that begins with a slash (/) as local regardless of which name service autofs uses.

Autofs Reference

This section describes more advanced autofs features and topics.

Autofs and Metacharacters

Autofs recognizes some characters as having a special meaning. For example, some characters are used for substitutions and some characters are used to protect other characters from the autofs map parser.

Ampersand (&)

If you have a map with many subdirectories specified, as in the following example, consider using string substitutions.

```
john      willow:/home/john
mary      willow:/home/mary
joe       willow:/home/joe
able      pine:/export/able
baker     peach:/export/baker
```

You can use the ampersand character (&) to substitute the key wherever the key appears. If you use the ampersand, the previous map changes to the following text:

```
john      willow:/home/&
mary      willow:/home/&
```

```
joe      willow:/home/&
able     pine:/export/&
baker    peach:/export/&
```

You could also use key substitutions in a direct map in situations such as the following example:

```
/usr/man    willow,cedar,poplar:/usr/man
```

You can also simplify the entry further as follows:

```
/usr/man    willow,cedar,poplar:&
```

Notice that the ampersand substitution uses the whole key string. Therefore, if the key in a direct map starts with a / (as it should), the slash is included in the substitution. Consequently, for example, you could not include the following entry:

```
/progs     &1,&2,&3:/export/src/progs
```

Autofs would interpret the example as follows:

```
/progs     /progs1,/progs2,/progs3:/export/src/progs
```

Asterisk (*)

You can use the universal substitute character, the asterisk (*), to match any key. For example, you could mount the /export file system from all hosts through this map entry.

```
*          &:/export
```

Each ampersand is substituted by the value of any given key. Autofs interprets the asterisk as an end-of-file character.

Autofs and Special Characters

If you have a map entry that contains special characters, you might have to mount directories that have names that the autofs map parser cannot process properly. The autofs parser is sensitive to names that contain colons, commas, and spaces, for example. These names should be enclosed in double quotes, as in the following example:

```
/vms      -ro    vmserver: - - - "rc0:dk1 - "  
/mac      -ro    gator:/ - "Mr Disk - "
```

Administering Network File Systems

This chapter provides information about how to perform NFS administration tasks such as setting up NFS services, adding new file systems to a share, and mounting file systems. This chapter also contains procedures that are used to configure and maintain NFS and FedFS referrals.

This chapter contains the following topics:

- “Automatic File System Sharing” on page 72
- “Mounting File Systems” on page 74
- “Setting Up the NFS Service” on page 81
- “Administering the Secure NFS System” on page 85
- “Administering WebNFS” on page 87
- “Administering NFS Referrals” on page 90
- “Administering FedFS” on page 91

Note - If your system has zones enabled and you want to use this feature in a non-global zone, see [Introduction to Oracle Solaris Zones](#).

About Administering Network File Systems

Your responsibilities as an NFS administrator depend on your site's requirements and your role as a network administrator. If you are responsible for all the systems on your local network, you might be responsible for determining the following:

- Which systems can serve as dedicated NFS servers
- Which systems can serve as both NFS servers and NFS clients
- Which systems serve as NFS clients only

Maintaining a server after it has been set up involves the following tasks:

- Sharing and unsharing file systems as necessary

- Modifying administrative files to update the lists of file systems your system mounts automatically
- Checking the status of the network
- Diagnosing and fixing NFS-related problems as they arise
- Setting up maps for autofs

A system can be both an NFS server and an NFS client. So, a system can be used to share local file systems with remote systems and to mount remote file systems.

Automatic File System Sharing

In the Oracle Solaris 11 release, the `share` command creates permanent shares that are automatically shared during system startup. Unlike previous releases, you do not need to edit the `/etc/dfs/dfstab` file to record the information about shares for subsequent reboots. This file is no longer used.

Task Map for File System Sharing

The following task map links to procedures that describe file system sharing using the NFS service.

TABLE 2 File System Sharing (Task Map)

| Task | Description | For Instructions |
|---|--|--|
| Establish automatic file system sharing | Configures a server so that file systems are automatically shared when the server is rebooted. | “How to Set Up Automatic File System Sharing” on page 72 |
| Enable NFS server logging | Configures a server so that NFS logging is run on selected file systems. | “How to Enable NFS Server Logging” on page 73 |

▼ How to Set Up Automatic File System Sharing

1. Become an administrator.

For more information, see [“Using Your Assigned Administrative Rights” in *Securing Users and Processes in Oracle Solaris 11.3*](#).

2. Define the file systems to be shared.

Note - NFS server processes do not start by default.

You can set the `share.nfs` property to share a ZFS file system.

```
$ zfs set share.nfs=on | off filesystem
```

For example:

```
$ zfs set share.nfs=on pond/amy
```

For information about the available command options, see the [zfs_share\(1M\)](#) man page.

3. Verify that the file system is shared.

You can use the `share` command to get a list of all the shared file systems.

For example:

```
$ share
pond_amy          /pond/amy          nfs sec=sys,rw
```

Next Steps The next step is to set up your `autofs` maps so that clients can access the file systems that you have shared on the server. For more information about setting up `autofs` maps, see [Table 6, “Tasks for Administering Autofs,”](#) on page 95.

▼ How to Enable NFS Server Logging

1. Become an administrator.

For more information, see [“Using Your Assigned Administrative Rights”](#) in *Securing Users and Processes in Oracle Solaris 11.3*.

2. Change file system configuration values.

You can change the settings in the `/etc/nfs/nfslog.conf` file in one of the following ways:

- Change the data associated with the `global` tag to edit the default settings for all file systems.
- Add a new tag for the file system.

For information about the format of the `/etc/nfs/nfslog.conf` file, see the [nfslog.conf\(4\)](#) man page.

3. Define the file systems to use NFS server logging.

Use the `share` command to define each file system. The tag that is used with the `log=tag` option must be specified in the `/etc/nfs/nfslog.conf` file.

The following example uses the default settings in the `global` tag.

```
$ share -F nfs -o ro,log=global /export/ftp
```

4. Verify that the options you specified are listed.

For example:

```
$ share -F nfs
export_share_man      /export/share/man   sec=sys,ro
usr_share_src         /usr/src            sec=sys,rw=eng
export_ftp            /export/ftp         public,log=global,sec=sys,ro
```

5. Verify that the NFS log daemon, `nfslogd`, is running.

```
$ ps -ef | grep nfslogd
```

6. Check the status of the `nfslogd` daemon.

```
$ svcadm restart network/nfs/server:default
```

Mounting File Systems

File systems can be mounted automatically when the system is booted, on demand from the command line, or through the automounter. The automounter provides many advantages over mounting at boot time or from the command line. However, many situations require a combination of all three methods. Additionally, several ways of enabling or disabling processes exist, depending on the options you use when mounting a file system.

Mounting File Systems Task Map

The following table lists the tasks that are associated with file system mounting.

TABLE 3 Mounting File Systems (Task Map)

| Task | Description | For Instructions |
|----------------------------------|--|--|
| Mount a file system at boot time | Enables a file system to be mounted whenever a system is rebooted. | “How to Mount a File System at Boot Time” on page 75 |

| Task | Description | For Instructions |
|--|---|--|
| Mount a file system by using a command | Mounts a file system when a system is running. This procedure is useful for testing. | “How to Mount a File System From the Command Line” on page 76 |
| Mount a file system with the automounter | Enables access to a file system on demand without using the command line. | “Mounting With the Automounter” on page 77 |
| Mount all file systems with mirror mounts | Mounts all of the file systems from one server. | “How to Mount All File Systems From a Server” on page 77 |
| Start client-side failover | Enables the automatic failover to a working file system if a server fails. | “How to Use Client-Side Failover” on page 78 |
| Disable mount access for a client | Disables the ability of one client to access a remote file system. | “How to Disable Mount Access for One Client” on page 78 |
| Provide access to a file system through a firewall | Enables access to a file system through a firewall by using the WebNFS protocol. | “How to Mount an NFS File System Through a Firewall” on page 79 |
| Mount a file system by using an NFS URL | Enables access to a file system by using an NFS URL. This process provides file system access without using the MOUNT protocol. | “How to Mount an NFS File System by Using an NFS URL” on page 79 |

▼ How to Mount a File System at Boot Time

This procedure shows how to mount file systems at boot time instead of using the autofs maps. This procedure must be completed on every NFS client that requires access to remote file systems.

1. Become an administrator.

For more information, see [“Using Your Assigned Administrative Rights” in *Securing Users and Processes in Oracle Solaris 11.3*](#).

2. Add an entry for the file system to the `/etc/vfstab` file.

Entries in the `/etc/vfstab` file have the following syntax:

```
#device      device      mount      FS      fsck      mount      mount
#to mount    to fsck     point      type    pass     at boot    options
```

For information about the `/etc/vfstab` file entries, see the [`vfstab\(4\)` man page](#).



Caution - NFS servers that also have NFS client `vfstab` entries must always specify the `bg` option to avoid a system hang during reboot. For more information, see the [`mount\(1M\)` man page](#).

3. Enable the NFS client service.

```
$ svcadm enable -r nfs/client
```

Note - The `nfs/client` service does not need to be enabled for manual mounts or when using `autofs` maps. The sole purpose of the `nfs/client` service is to mount the NFS file systems that are listed in `/etc/vfstab` with the `mount at boot` field set to `yes`.

Example 3 Entry in the Client's `/etc/vfstab` File

Assume that you want a client system to mount the `/var/mail` directory from the server `wasp`. You want the file system to be mounted as `/var/mail` on the client, and you want the client to have read-write access. You would add the following entry to the client's `vfstab` file:

```
wasp:/var/mail - /var/mail nfs - yes rw
```

▼ How to Mount a File System From the Command Line

Mounting a file system from the command line is often performed to test a new mount point. This type of mount enables temporary access to a file system that is not available through the automounter. You can unmount the file system with the `umount` command or by rebooting the local system.

1. Become an administrator.

For more information, see [“Using Your Assigned Administrative Rights” in *Securing Users and Processes in Oracle Solaris 11.3*](#).

2. Mount the file system.

```
mount -F nfs -o specific-options resource mount-point
```

For example:

```
$ mount -F nfs -o ro bee:/export/share/local /mnt
```

For more information, see the [`mount_nfs\(1M\)`](#) man page.

In this example, the `/export/share/local` file system from the server `bee` is mounted read-only on `/mnt` on the local system.



Caution - The `mount` command does not provide warnings about invalid options. The command silently ignores any options that cannot be interpreted. To prevent unexpected behavior, verify all of the options that you use.

Mounting With the Automounter

Without any changes to the generic system, NFS clients can access remote file systems through the `/net` mount point. For information about establishing and supporting mounts with the automounter, see [Table 6, “Tasks for Administering Autofs,” on page 95](#). Type the following command to mount the `/export/share/local` file system:

```
$ cd /net/bee/export/share/local
```

Because the automounter enables all users to mount file systems, root access is not required. The automounter also automatically unmounts file systems, so you do not have to unmount file systems manually after you no longer need to access them.

▼ How to Mount All File Systems From a Server

The automatic mirror mount facility enables an NFS client to access all available file systems shared using NFS from an NFS server after one mount from that server has succeeded. The mirror mount occurs automatically and you only need to access the file system. For more information, see [“How Mirror Mounts Work” on page 49](#).

- 1. Become an administrator.**

For more information, see [“Using Your Assigned Administrative Rights” in *Securing Users and Processes in Oracle Solaris 11.3*](#).

- 2. Mount the root of the exported namespace of the server.**

This command mirrors the file system hierarchy from the server on the client. In this example, a `/mnt/export/share/local` directory structure is created.

```
$ mount bee:/ /mnt
```

- 3. Access a file system.**

This command or any other command that accesses the file system causes the file system to be mounted.

```
$ cd /mnt/export/share/local
```

▼ How to Use Client-Side Failover

1. Become an administrator.

For more information, see [“Using Your Assigned Administrative Rights” in *Securing Users and Processes in Oracle Solaris 11.3*](#).

2. On the NFS client, mount the file system by using the `ro` option.

You can mount from the command line, through the automounter, or by adding an entry to the `/etc/vfstab` file that resembles the following:

```
bee,wasp:/export/share/local - /usr/local nfs - no ro
```

Note - Servers that are running different versions of the NFS protocol cannot be mixed in a command issued at the command line or in a `vfstab` entry. Mixing servers that support NFS Version 2, NFS Version 3, or NFS Version 4 protocols can be performed only with `autofs`. In `autofs`, the best subset of NFS Version 2, NFS Version 3, or NFS Version 4 servers is used.

▼ How to Disable Mount Access for One Client

1. Become an administrator.

For more information, see [“Using Your Assigned Administrative Rights” in *Securing Users and Processes in Oracle Solaris 11.3*](#).

2. Disable mount access for one client.

```
$ share -F nfs -o specific_options pathname
```

For example:

```
$ share -F nfs ro=-rose:eng /export/share/man
```

```
ro=-rose:eng      Access list that allows read-only mount access to all clients in the eng
                  network group except for the host named rose
```

```
/export/share/   File system to be shared
man
```

▼ How to Mount an NFS File System Through a Firewall

Before You Begin This procedure requires that the file system on the NFS server be shared by using the `public` option. Additionally, any firewalls between the NFS client and the NFS server must allow TCP connections on port 2049. All file systems that are shared allow for public file handle access, so the `public` option is applied by default.

1. **Become an administrator.**
For more information, see [“Using Your Assigned Administrative Rights” in *Securing Users and Processes in Oracle Solaris 11.3*](#).
2. **Manually mount the file system by using a command such as the following:**

```
$ mount -F nfs host:pathname mount-point
```

For example:

```
$ mount -F nfs bee:/export/share/local /mnt
```

In this example, the file system `/export/share/local` is mounted on the local client by using the public file handle. An NFS URL can be used instead of the standard path name. If the public file handle is not supported by the server `bee`, the mount operation fails.

Mount an NFS File System by Using an NFS URL

You can choose to include the `public` option with an NFS URL. Without the `public` option, the MOUNT protocol is used if the public file handle is not supported by the server. The `public` option forces the use of the public file handle, and the mount fails if the public file handle is not supported.

Note - The NFS protocol version that is used when you mount the file system is the highest version supported by both the NFS client and the NFS server. However, you can use the `vers=#` option to select a specific NFS protocol version.

▼ How to Mount an NFS File System by Using an NFS URL

1. **Become an administrator.**
For more information, see [“Using Your Assigned Administrative Rights” in *Securing Users and Processes in Oracle Solaris 11.3*](#).

2. Manually mount the file system by using an NFS URL.

```
$ mount -F nfs nfs://host[:port]/pathname mount-point
```

Example 4 Mounting an NFS File System by Using an NFS URL

```
$ mount -F nfs nfs://bee:3000/export/share/local /mnt
```

In this example, the `/export/share/local` file system is being mounted from the server `bee` by using NFS port number `3000`. The port number is not required, and by default the standard NFS port number of `2049` is used.

Displaying Information About File Systems Available for Mounting

The `showmount` command displays information about file systems that have been remotely mounted or are available for mounting. You use the `-e` option to list the shared file systems. For example:

```
$ /usr/sbin/showmount -e bee
export list for bee:
/export/share/local  (everyone)
/export/home         tulip,lilac
/export/home2       rose
```

For information about other options, see the [showmount\(1M\)](#) man page.

In some environments, information about shared file systems and the systems that have mounted them should not be displayed. You can set the `showmount_info` property of the `sharectl` command to `none`, which ensures that the client cannot view the following file system information:

- Information about file systems that the client cannot access
- Information about all the shared file systems
- Information about other systems that have mounted the file systems

EXAMPLE 5 Restricting File System Information Displayed to Clients

```
bee$ sharectl set -p showmount_info=none nfs
```

The following information is displayed on the client `rose`:


```
$ /usr/sbin/showmount -e bee
export list for bee:
/export/share/local (everyone)
/export/home2      rose
```

The information about the `/export/home` file system is no longer displayed.

Setting Up the NFS Service

This section describes some of the tasks that are necessary to set up the NFS service.

Note - NFS Version 4 is the default version of NFS supported in Oracle Solaris 11.2 and later.

TABLE 4 Setting Up the NFS Service

| Task | Description | For Instructions |
|----------------------------------|--|--|
| Start and stop the NFS server | Starts the NFS service if it has not been started automatically. Stops the NFS service. Normally, the service does not need to be stopped. | “Starting and Stopping the NFS Service” on page 81 |
| Start and stop the automounter | Starts and stops the automounter. This procedure is required when some of the automounter maps are changed. | “Starting and Stopping the Automounter” on page 81 |
| Select different versions of NFS | Selects a NFS version other than NFS Version 4 on the server and clients. | “Selecting Different Versions of NFS” on page 82 |

Starting and Stopping the NFS Service

As an administrator, use the `svcadm` command to enable and disable the NFS service on the server.

- To enable the NFS service on the server:

```
$ svcadm enable network/nfs/server
```

- To disable the NFS service on the server:

```
$ svcadm disable network/nfs/server
```

Starting and Stopping the Automounter

As an administrator, use the `svcadm` command to enable and disable the `autofs` daemon.

- To enable the autofs daemon:

```
$ svcadm enable system/filesystem/autofs
```

- To disable the autofs daemon:

```
$ svcadm disable system/filesystem/autofs
```

Selecting Different Versions of NFS

If you want to use a version of NFS other than NFS Version 4, you can select a different version:

- If you want to select a different version of NFS on the NFS server, see [“How to Select Different Versions of NFS on a Server” on page 82](#).
- If you want to select a different version of NFS on the NFS client, see [“How to Select Different Versions of NFS on a Client” on page 83](#).
- If you want to select a different version of NFS on the NFS client using the command line, see [“How to Use the mount Command to Select Different Versions of NFS on a Client” on page 84](#).

▼ How to Select Different Versions of NFS on a Server

You can select another version of NFS if you choose not to use NFS Version 4, which is set by default.

- 1. Become an administrator.**

For more information, see [“Using Your Assigned Administrative Rights” in *Securing Users and Processes in Oracle Solaris 11.3*](#).

- 2. Change SMF parameters to set the NFS Version numbers.**

For example, if you want the server to provide only NFS Version 3, set the values for both the `server_versmax` and `server_versmin` parameters to 3.

```
$ sharectl set -p server_versmax=3 nfs
$ sharectl set -p server_versmin=3 nfs
```

- 3. If you want to disable server delegation, change the `server_delegation` property.**

```
$ sharectl set -p server_delegation=off nfs
```

NFS server delegation enables an NFS client to cache files until another NFS client needs access to the same files. In NFS Version 4, server delegation is enabled by default. For more information, see [“Delegation in NFS Version 4” on page 35](#).

4. If you want to set a common domain for clients and servers, change the `nfsmapid_domain` property.

You can set a common domain for the clients and servers to enable user ID or group ID mapping between the client and the server.

```
$ sharectl set -p nfsmapid_domain=my.example.com nfs
```

where *my.example.com* provides the common domain name.

For more information about the `nfsmapid` daemon, see [“NFS Daemons” on page 160](#).

5. Check whether the NFS service is running on the server.

```
$ svcs network/nfs/server
```

6. If necessary, enable the NFS service.

If the NFS service is offline, type the following command to enable the service:

```
$ svcadm enable network/nfs/server
```

For information about configuring the NFS service, see [“How to Set Up Automatic File System Sharing” on page 72](#).

See Also [“Version Negotiation in NFS” on page 28](#)

▼ How to Select Different Versions of NFS on a Client

The following procedure explains how to control which version of NFS is used on the client. The NFS version that is set by default is NFS Version 4.

1. Become an administrator.

For more information, see [“Using Your Assigned Administrative Rights” in *Securing Users and Processes in Oracle Solaris 11.3*](#).

2. Change SMF parameters to set the NFS Version numbers.

For example, if you want all file systems to be mounted using the NFS Version 3 protocol, set the values for both the `client_versmax` and `client_versmin` parameters to 3.

```
$ sharectl set -p client_versmax=3 nfs
```

```
$ sharectl set -p client_versmin=3 nfs
```

3. Mount NFS on the client.

```
$ mount server-name:/share-point /local-dir
```

| | |
|---------------------|-------------------------------|
| <i>server-name</i> | Name of the server. |
| <i>/share-point</i> | Path of the remote directory |
| <i>/local-dir</i> | Path of the local mount point |

See Also [“Version Negotiation in NFS” on page 28](#)

▼ How to Use the `mount` Command to Select Different Versions of NFS on a Client

This procedure explains how to use the `mount` command to control which version of NFS is used on a client for a particular mount. To find out how to modify the NFS version for all file systems mounted by the client, see [“How to Select Different Versions of NFS on a Client” on page 83](#).

1. Become an administrator.

For more information, see [“Using Your Assigned Administrative Rights” in *Securing Users and Processes in Oracle Solaris 11.3*](#).

2. Mount the desired version of NFS on the client.

```
$ mount -o vers=value server-name:/share-point /local-dir
```

| | |
|---------------------|-------------------------------|
| <i>value</i> | NFS version number |
| <i>server-name</i> | Name of the server |
| <i>/share-point</i> | Path of the remote directory |
| <i>/local-dir</i> | Path of the local mount point |

Note - This command overrides the NFS client settings in the SMF repository.

See Also [“Version Negotiation in NFS” on page 28](#)

Administering the Secure NFS System

To use the Secure NFS system, all the systems that you are responsible for must have a domain name. Typically, a domain is an administrative entity of several systems that is part of a larger network. If you are running a name service, you should also establish the name service for the domain. For information about name services, see [Working With Oracle Solaris 11.3 Directory and Naming Services: DNS and NIS](#).

Kerberos V5 authentication is supported by the NFS service. For more information, see [Chapter 2, “Kerberos on Oracle Solaris” in Managing Kerberos and Other Authentication Services in Oracle Solaris 11.3](#).

You can also configure the Secure NFS environment to use Diffie-Hellman authentication. For more information about Diffie-Hellman authentication, see [Chapter 9, “Configuring Network Services Authentication” in Managing Kerberos and Other Authentication Services in Oracle Solaris 11.3](#).

▼ How to Set Up a Secure NFS Environment With DH Authentication

1. Assign a domain name.

Make the domain name known to each system in the domain. For information about setting up a machine's NIS domain name, see [“How to Set a Machine's NIS Domain Name” in Working With Oracle Solaris 11.3 Directory and Naming Services: DNS and NIS](#).

```
$ domainname domain-name
```

2. Establish public keys and secret keys for your clients' users by using the `newkey` command.

```
$ newkey -u username -s name-service
```

Users can establish personal secure RPC passwords by using the `chkey` command.

```
$ chkey -p -s name-service -m mechanism
```

When public keys and secret keys have been generated, the public keys and encrypted secret keys are stored in the `publickey` database.

For information about these commands, see the [`newkey\(1M\)`](#) and the [`chkey\(1\)`](#) man pages.

3. Verify that the name service is responding.

For example:

- If you are running NIS, verify that the ypbind daemon is running. For more information, see [“ypbind Not Running on NIS Client”](#) in *Working With Oracle Solaris 11.3 Directory and Naming Services: DNS and NIS*.
- If you are running LDAP, verify that the ldap_cachemgr daemon is running. For more information, see [“Monitoring LDAP Client Status”](#) in *Working With Oracle Solaris 11.3 Directory and Naming Services: LDAP*.

4. Verify that the key serv daemon of the key server is running.

```
$ ps -ef | grep key serv
root    100      1  16   Apr 11 ?           0:00 /usr/sbin/key serv
root    2215     2211  5 09:57:28 pts/0    0:00 grep key serv
```

If the daemon is not running, type the following to start the key server:

```
$ svcadm enable network/rpc/key serv
```

5. Decrypt and store the secret key.

Usually, the login password is identical to the network password. In this situation, keylogin is not required. If the passwords are different, the users have to log in, and then run keylogin. You still need to use the keylogin -r command as root to store the decrypted secret key in /etc/.rootkey.

Note - You need to run keylogin -r if the root secret key changes or if the /etc/.rootkey file is lost.

6. Set the security mode for the file system to be shared.

For Diffie-Hellman authentication add the sec=dh option to the command line.

```
$ share -F nfs -o sec=dh /export/home
```

For more information about security modes, see the [nfssec\(5\)](#) man page.

7. Update the automounter maps for the file system.

If you are using Diffie-Hellman authentication, edit the auto_master data to include sec=dh as a mount option in the appropriate entries.

```
/home auto_home -nosuid,sec=dh
```

When you reinstall, move, or upgrade a system, remember to save the `/etc/.rootkey` file if you do not establish new keys or change the keys for root. If you delete the `/etc/.rootkey` file, type the following command:

```
$ keylogin -r
```

Administering WebNFS

This section provides instructions for administering the WebNFS system.

TABLE 5 Administering WebNFS (Task Map)

| Task | Description | For Instructions |
|--|--|--|
| Plan for WebNFS | Issues to consider before enabling the WebNFS service. | “Planning for WebNFS Access” on page 87 |
| Enable WebNFS | Enables mounting of an NFS file system by using the WebNFS protocol. | “How to Enable WebNFS Access” on page 88 |
| Enable WebNFS through a firewall | Enables access to files through a firewall by using the WebNFS protocol. | “Enabling WebNFS Access Through a Firewall” on page 90 |
| Browse by using an NFS URL | Uses an NFS URL within a web browser. | “Accessing an NFS URL by Using a Browser” on page 89 |
| Use a public file handle with autofs | Uses public file handle when mounting a file system with the automounter. | “How to Use a Public File Handle With Autofs” on page 107 |
| Use an NFS URL with autofs | Adds an NFS URL to the automounter maps. | “How to Use NFS URLs With Autofs” on page 107 |
| Provide access to a file system through a firewall | Enables access to a file system through a firewall by using the WebNFS protocol. | “How to Mount an NFS File System Through a Firewall” on page 79 |
| Mount a file system by using an NFS URL | Enables access to a file system by using an NFS URL. This process enables file system access without using the MOUNT protocol. | “How to Mount an NFS File System by Using an NFS URL” on page 79 |

Planning for WebNFS Access

To use WebNFS, you first need an application that is capable of running and loading an NFS URL (for example, `nfs://server/path`). The next step is to choose the file system that can be exported for WebNFS access. If the application is web browsing, often the document root

for the web server is used. You need to consider several factors when choosing a file system to export for WebNFS access.

- Each server has one public file handle that by default is associated with the server's root file system. The path in an NFS URL is evaluated relative to the directory with which the public file handle is associated. If the path leads to a file or directory within an exported file system, the server provides access. You can use the `public` option of the `share` command to associate the public file handle with a specific exported directory. Using this option enables URLs to be relative to the shared file system rather than to the server's root file system. The root file system does not allow web access unless the root file system is shared.
- The WebNFS environment enables users who already have mount privileges to access files through a browser. This capability is enabled regardless of whether the file system is exported by using the `public` option. Because users already have access to these files through the NFS setup, this access should not create any additional security risk. You only need to share a file system by using the `public` option if users who cannot mount the file system need to use WebNFS access.
- File systems that are already open to the public make good candidates for using the `public` option. Some examples are the top directory in an ftp archive or the main URL directory for a web site.
- You can use the `index` option with the `share` command to force the loading of an HTML file. Otherwise, you can list the directory when an NFS URL is accessed.

After a file system is chosen, review the files and set access permissions to restrict viewing of files or directories, as needed. Establish the permissions, as appropriate, for any NFS file system that is being shared. For many sites, 755 permissions for directories and 644 permissions for files provide the correct level of access.

You need to consider additional factors if both NFS and HTTP URLs are to be used to access one web site. For more information about WebNFS limitations, see [“WebNFS Limitations With Web Browser Use” on page 45](#).

▼ How to Enable WebNFS Access

Before You Begin By default, all file systems that are available for NFS mounting are automatically available for WebNFS access. Use this procedure for one of the following reasons:

- To allow NFS mounting on a server that does not currently allow NFS mounting
- To reset the public file handle to shorten NFS URLs by using the `public` option with the `share` command
- To force an HTML file to be loaded by using the `index` option with the `share` command

You can also use the `sharectl` utility to configure file-sharing protocols such as NFS. For more information about configuring file sharing protocols, see the [sharectl\(1M\)](#) man page.

For information about issues to consider before starting the WebNFS service, see “[Planning for WebNFS Access](#)” on page 87.

1. Become an administrator.

For more information, see “[Using Your Assigned Administrative Rights](#)” in *Securing Users and Processes in Oracle Solaris 11.3*.

2. Define the file systems to be shared by the WebNFS service.

Use the `share` command to define each file system.

```
$ share -F nfs -o specific-options pathname
```

For information about the available options for the `share_nfs` command, see the [share_nfs\(1M\)](#) man page.

3. Verify that the options you specified are listed.

```
$ share -F nfs
```

For example:

```
$ share -F nfs
export_share_man  /export/share/man  sec=sys,ro
usr_share_src    /usr/src      sec=sys,rw=eng
export_ftp       /export/ftp   sec=sys,ro,public,index=index.html
```

Accessing an NFS URL by Using a Browser

Browsers that are capable of supporting the WebNFS service should provide access to an NFS URL that resembles the following:

```
nfs://server[:port]/path
```

| | |
|---------------|--|
| <i>server</i> | Name of the file server |
| <i>port</i> | Port number to use (2049, default value) |
| <i>path</i> | Path to file, which can be relative to the public file handle or to the root file system |

Note - In most browsers, the URL service type (for example, `nfs` or `http`) is remembered from one transaction to the next. The exception occurs when a URL that includes a different service type is loaded. For example, if a reference to an HTTP URL is loaded after you use an NFS URL, subsequent pages are loaded by using the HTTP protocol instead of the NFS protocol.

Enabling WebNFS Access Through a Firewall

You can enable WebNFS access for clients that are not part of the local subnet by configuring the firewall to allow a TCP connection on port 2049. Just allowing access for `httpd` does not allow NFS URLs to be used.

Administering NFS Referrals

An NFS referral enables an NFS Version 4 server to point to file systems that are located on other NFS Version 4 servers as a way of connecting multiple NFS Version 4 servers into a uniform namespace.

▼ How to Create and Access an NFS Referral

1. On an NFS server, create a referral.

Add the referral on an NFS-shared file system, pointing to one or more existing NFS-shared file systems. For example:

```
server1 nfsref add /share/docs server2:/usr/local/docs server3:/tank/docs
Created reparse point /share/docs
```

2. Verify that the referral was created.

```
server1$ nfsref lookup /share/docs
/share/docs points to:
server2:/usr/local/docs
server3:/tank/docs
```

3. On the client, access the mount point to mount the referral.

```
client1$ ls /share/docs
```

If the mount fails, check the connectivity on the NFS client and check the shared file system on the NFS server. For more information about troubleshooting NFS, see [“NFS Troubleshooting Procedures” on page 140](#).

Example 6 Modifying an Existing NFS Referral

To add another file system, such as `server4:/tank/docs`, to the existing referral that was created in this procedure, you would type the command from Step 2 with the new file system.

```
server1$ nfsref add /share/docs server2:/usr/local/docs \
server3:/tank/docs server4:/tank/docs
```

The `add` subcommand replaces the information in the current referral with the new information from the command.

▼ How to Remove an NFS Referral

- To remove an NFS referral, type the following command:

```
server1$ nfsref remove /share/docs
Removed svc_type 'nfs-basic' from /share/docs
```

This removes a single referral that was created at `/share/docs`.

Administering FedFS

You use the FedFS protocol to construct and maintain a federated file system. This file system can include many different file servers to create a multivendor global namespace.

Set Up a DNS Record for a FedFS Server

After an appropriate DNS record is created, mounting a file system using FedFS is completed by the automounter after the mount point has been accessed. The DNS record for the server appears similar to the following:

```
$ nslookup -q=svr _nfs-domainroot._tcp.example.com bee.example.com
```

```
Server:      bee.example.com
Address:     192.0.2.0
```

```
_nfs-domainroot._tcp.example.com      service = 1 0 2049 bee.example.com.
```

After you set up the DNS record, Oracle Solaris will mount the FedFS file system automatically when an application accesses the `/nfs4/example.com` mount point.

▼ How to Create a Namespace Database

A namespace database (NSDB) is used to provide information about the set of files from different types of servers that are combined into a single FedFS namespace. This procedure is performed on the LDAP server.

Before You Begin You must have an LDAP server installed. For more information, see [Chapter 4, “Setting Up the Oracle Directory Server Enterprise Edition With LDAP Clients”](#) in *Working With Oracle Solaris 11.3 Directory and Naming Services: LDAP*.

- 1. Become an administrator.**

For more information, see [“Using Your Assigned Administrative Rights”](#) in *Securing Users and Processes in Oracle Solaris 11.3*.

- 2. Update the `/etc/openldap/slapd.conf` file with the following entries:**

```
include      /usr/lib/fs/nfs/fedfs-11.schema
suffix dc=example,dc=org
rootdn cn=Manager,dc=example,dc=org
rootpw password
```

- 3. Create a unique name for the FedFS data.**

```
$ nsdb-update-nci -l NSDB -r port -D bind_DN -w bind-PW nce
```

For example:

```
$ nsdb-update-nci -l localhost -r 389 -D cn=Manager -w\
  example.org dc=example,dc=org adding new entry "dc=example,dc=org"
NCE entry created
```

`-l` Specifies the LDAP server implementing the NSDB

`-r` Specifies the port on which the LDAP server implementing the NSDB is listening

- D Specifies the unique name of a user permitted to change the NSDB information
- w Specifies the password for the bind DN user

For more information, see the [nsdb-update-nci\(1M\)](#) man page.

▼ How to Use a Secured Connection to the NSDB

Before You Begin You must have an LDAP server installed. For more information, see [Chapter 4, “Setting Up the Oracle Directory Server Enterprise Edition With LDAP Clients”](#) in *Working With Oracle Solaris 11.3 Directory and Naming Services: LDAP*.

1. Become an administrator.

For more information, see [“Using Your Assigned Administrative Rights”](#) in *Securing Users and Processes in Oracle Solaris 11.3*.

2. On the LDAP server, create a certificate.

You need a certificate to secure the LDAP traffic.

```
$ mkdir /etc/openldap/certs
$ mkdir /etc/openldap/certs/keys
$ cd /etc/openldap/certs
$ openssl req -x509 -nodes -days 3650 -newkey rsa:2048 \
  -keyout keys/ldapskey.pem -out ldapscert.pem
$ chown -R openldap:openldap /etc/openldap/certs/*
$ chmod 0400 keys/ldapskey.pem
```

3. Add declarations to the `/etc/openldap/slapd.conf` file.

```
TLSCertificateFile /etc/openldap/certs/ldapscert.pem
TLSCertificateKeyFile /etc/openldap/certs/keys/ldapskey.pem
```

4. Copy the certificate to the NFS server and clients.

```
$ scp ldap-server:/etc/openldap/certs/keys/ldapskey.pem \
  /etc/openldap/certs/keys/ldapskey.pem
$ chmod 0400 /etc/openldap/certs/keys/ldapskey.pem
```

5. On the NFS server and clients, update the connection entry.

```
$ nsdbparams update -f ldapscert.pem -t FEDFS_SEC_TLS localhost
```

For information about options available with the `nsdbparams` command, see the [nsdbparams\(1M\)](#) man page.

▼ How to Create a FedFS Referral

Before You Begin You must have an NFS server installed.

1. **Become an administrator.**
2. **Create a connection entry for an NSDB.**

This command creates a connection entry between the NSDB that is defined on the LDAP server and an NFS server.

```
$ nsdbparams update -D cn=Manager,dc=example,dc=org -w example.org nsdb.example.org
```

3. **Create a FedFS referral.**

```
$ nfsref -t svc-type add path location
```

`-t svc-type` Specifies the service type of the referral

For example:

```
$ nfsref -t nfs-fedfs add /share/docs server2:/usr/local/docs server3:/tank/docs  
Created reparse point /share/doc
```

Administering Autofs

This chapter provides information about how to perform autofs administration tasks such as accessing file systems, modifying autofs maps, and using security restrictions with autofs.

This chapter contains the following topics:

- [“Autofs Administration” on page 95](#)
- [“Using SMF Parameters to Configure Your Autofs Environment” on page 97](#)
- [“Modifying the Maps” on page 98](#)
- [“Customizing the Automounter” on page 100](#)

Note - If your system has zones enabled and you want to use this feature in a non-global zone, see [Introduction to Oracle Solaris Zones](#).

Autofs Administration

The following table provides a description and a pointer to many of the tasks that are related to autofs.

TABLE 6 Tasks for Administering Autofs

| Task | Description | For Instructions |
|---|--|--|
| Start and stop autofs | Start and stop the automount service without having to reboot the system | “Starting and Stopping the Automounter” on page 81 |
| Configure your autofs environment through the autofs SMF parameters | Assign values to parameters in the SMF repository | “Using SMF Parameters to Configure Your Autofs Environment” on page 97 |
| Access file systems by using autofs | Access file systems by using the automount service | “Mounting With the Automounter” on page 77 |
| Modify the autofs master map | The master map is used to list other maps | “How to Modify the Master Map” on page 98 |
| Modify an autofs indirect map | An indirect map is used for most maps | “How to Modify Indirect Maps” on page 99 |

| Task | Description | For Instructions |
|---|---|---|
| Modify an autofs direct map | A direct map is used to establish a direct association between a mount point on a client and a server | “How to Modify Direct Maps” on page 99 |
| Modify the autofs maps to access non-NFS file systems | Set up an autofs map with an entry for a CD-ROM application | “Accessing Non-NFS File Systems” on page 100 |
| Use /home maps | <ul style="list-style-type: none"> ■ Set up a common /home map ■ Set up a /home map that refers to multiple file systems | <ul style="list-style-type: none"> ■ “Setting Up a Common View of /home” on page 100 ■ “How to Set Up /home With Multiple Home Directory File Systems” on page 101 |
| Using a new autofs mount point | <ul style="list-style-type: none"> ■ Set up a project-related autofs map ■ Set up an autofs map that supports different client architectures ■ Set up an autofs map that supports different operating systems | <ul style="list-style-type: none"> ■ “How to Consolidate Project-Related Files Under a Common Directory” on page 102 ■ “How to Set Up Different Architectures to Access a Shared Namespace” on page 104 ■ “How to Support Incompatible Client Operating System Versions” on page 105 |
| Replicate file systems with autofs | Provide access to file systems that fail over | “How to Replicate Shared Files Across Several Servers” on page 106 |
| Using security restrictions with autofs | Provide access to file systems while restricting remote root access to the files | “How to Apply Autofs Security Restrictions” on page 106 |
| Using a public file handle with autofs | Force use of the public file handle when mounting a file system | “How to Use a Public File Handle With Autofs” on page 107 |
| Using an NFS URL with autofs | Add an NFS URL so that the automounter can use it | “How to Use NFS URLs With Autofs” on page 107 |
| Disable autofs browsability | <ul style="list-style-type: none"> ■ Disable browsability so that autofs mount points are not automatically populated on a single client ■ Disable browsability so that autofs mount points are not automatically populated on all clients ■ Disable browsability so that a specific autofs mount point is not automatically populated on a client | <ul style="list-style-type: none"> ■ “How to Completely Disable Autofs Browsability on a Single NFS Client” on page 108 ■ “How to Disable Autofs Browsability for All Clients” on page 108 ■ “How to Disable Autofs Browsability on a Selected File System” on page 108 |

Using SMF Parameters to Configure Your Autofs Environment

You can use SMF parameters to configure your autofs environment. Specifically, this facility provides an additional way to configure your autofs commands and autofs daemons. You can make the same specifications with the `sharectl` command that you would make on the command line. You can make your specifications by providing values to keywords.

▼ How to Configure Your Autofs Environment Using SMF Parameters

1. Become an administrator.

For more information, see [“Using Your Assigned Administrative Rights” in *Securing Users and Processes in Oracle Solaris 11.3*](#).

2. Add or modify an autofs SMF parameter.

For example, if you want to turn off browsing for all autofs mount points, use the following command:

```
$ sharectl set -p nobrowse=on autofs
```

The `nobrowse` keyword is equivalent to the `-n` option of the `automountd` command. For information about parameters supported for autofs, see [autofs\(4\)](#).

3. Restart the autofs daemon.

```
$ svcadm restart system/filesystem/autofs
```

Administrative Tasks Involving Maps

Your choice of a map and name service type affects the mechanism that you need to use to make changes to the autofs maps.

Note - Use indirect maps whenever possible. Indirect maps are easier to construct and less demanding on the systems' file systems. Also, indirect maps do not occupy as much space in the mount table as direct maps.

The types of maps and their uses are:

- Master – Associates a directory with a map
- Direct – Directs autofs to specific file systems
- Indirect – Directs autofs to reference-oriented file systems

The way that you make changes to your autofs environment depends on your name service. To make changes if you are using local files as a name service, use a text editor. If your name service is NIS or LDAP, make changes as required in these name services.

You might have to run the automount command depending on the modification you have made to the type of map. For example, if you have made an addition or a deletion to a direct map, you need to run the automount command on the local system. By running the command, you make the change effective. However, if you have modified an existing entry, you do not need to run the automount command for the change to become effective. You always have to run the automount command if you make changes to the master map. You never have to run the automount command if you make changes to the indirect map.

Modifying the Maps

This section describes how to update several types of automounter maps.

▼ How to Modify the Master Map

1. **Log in as a user who has permissions to change the maps based on the name service that you are using. If you are using the local map files, assume the root role.**

2. **Make your changes to the master map.**

The specific steps needed to change the map depends on the name service that you are using. If you are using local files as a name service, use a text editor. If your name service is NIS, use make files.

3. **For each client, become an administrator.**

For more information, see [“Using Your Assigned Administrative Rights” in *Securing Users and Processes in Oracle Solaris 11.3*](#).

4. **For each client, run the automount command to make your changes effective.**

5. **Notify your users to run the `automount` command as superuser on their own systems in order to incorporate the new information from the master map.**

▼ How to Modify Indirect Maps

1. **Log in as a user who has permissions to change the maps.**
2. **Make your changes to the indirect map.**

The specific steps needed to change the map depends on the name service that you are using.

▼ How to Modify Direct Maps

1. **Log in as a user who has permissions to change the maps.**
2. **Make your changes to the direct map.**

The specific steps needed to change the map depends on the name service that you are using.

3. **Notify your users of the changes.**

Notification is required so that the users can run the `automount` command as superuser on their own systems, if necessary.

Note - If you only modify or change the contents of an existing direct map entry, you do not need to run the `automount` command.

For example, suppose you modify the `auto_direct` map so that the `/usr/src` directory is now mounted from a different server. If `/usr/src` is not mounted at this time, the new entry becomes effective immediately when you try to access `/usr/src`. If `/usr/src` is mounted now, you can wait until the auto-unmounting occurs, then access the file.

Avoiding Mount Point Conflicts

If you have a local disk partition that is mounted on `/src` and you plan to use the `autofs` service to mount other source directories, the NFS service hides the local partition whenever you try to reach it. Therefore, You need to mount the partition in some other location.

For example, to mount the partition on `/export/src` you would add an entry in the `/etc/vfstab` file such as the following:

```
/dev/dsk/d0t3d0s5 /dev/rdisk/c0t3d0s5 /export/src ufs 3 yes -
```

You also need to add an entry in `auto_src`. In this example, the name of the system is `terra`.

```
terra terra:/export/src
```

Accessing Non-NFS File Systems

Autofs can also mount files other than NFS files, for example, files on removable media, such as CD-ROM or USB flash drives.

Instead of mounting a file system from a server, you put the media in the drive and reference the file system from the map. For example, to access a CD-ROM application, become an administrator and add an entry for the CD-ROM file system similar to the following example in the `autofs` map, with the CD-ROM device name following the colon:

```
hsfs -fstype=hsfs,ro :/dev/sr0
```

Customizing the Automounter

This section describes how to customize the automounter maps to provide an easy-to-use directory structure.

Setting Up a Common View of `/home`

The ideal is for all network users to be able to locate their own home directories or the home directories of other users under `/home`. This view should be common across all systems, whether client or server.

Every Oracle Solaris installation comes with a master map: `/etc/auto_master`.

```
# Master map for autofs
#
+auto_master
/net      -hosts      -nosuid,nobrowse
/home     auto_home   -nobrowse
/nfs4     -fedfs      -ro,nosuid,nobrowse
```

A map for `auto_home` is also installed under `/etc`. When a new local user is created, an entry is automatically added to `/etc/auto_home`. For example:

```
# Home directory map for autofs
#
rusty dragon:/export/home/&
+auto_home
```

On the server named `dragon`, the home directory for `rusty` can be accessed through `/export/home/rusty` as well as `/home/rusty`.

With the `auto_home` map in place, users can refer to any home directory (including their own) with the path `/home/user`. `user` is their login name and the key in the map. This common view of all home directories is valuable when logging in to another user's system. Autofs mounts your home directory for you. Similarly, if you run a remote windowing system client on another system, the client program has the same view of the `/home` directory. This common view also extends to the server.

Users do not need to be aware of the real location of their home directories. If a user needs more disk space and needs to have the home directory relocated to another server, you need only change the user's entry in the `auto_home` map to reflect the new location. Other users can continue to use the `/home/user` path.

Note - Do not permit users to run `setuid` executables from their home directories. Without this restriction, any user could have superuser privileges on any system.

▼ How to Set Up /home With Multiple Home Directory File Systems

1. Become an administrator.

For more information, see [“Using Your Assigned Administrative Rights” in *Securing Users and Processes in Oracle Solaris 11.3*](#).

2. Install home directory partitions under `/export/home`.

If the system has several partitions, install the partitions under separate directories, for example, `/export/home1` and `/export/home2`.

3. Update the `auto_home` map.

Whenever you create a new user account, type the location of the user's home directory in the `auto_home` map. Map entries can be simple, for example:

```
user1      system1:/export/home1/&
user2      system1:/export/home1/&
user3      system2:/export/home2/&
user4      system1:/export/home3/&
```

Notice the use of the & (ampersand) as a substitute for the map key. The ampersand is an abbreviation for the second occurrence of user1 in the following example:

```
user1      system1:/export/home1/user1
```

▼ How to Consolidate Project-Related Files Under a Common Directory

You can use `autofs` to consolidate files in a directory that is common across multiple systems. You can add the directory structure of the project-related files to the `autofs` map for the common directory. This directory structure enables the users to use the project files irrespective of physical and hardware changes in the systems.

- 1. Become an administrator.**

For more information, see [“Using Your Assigned Administrative Rights” in *Securing Users and Processes in Oracle Solaris 11.3*](#).

- 2. Add an entry for the common directory to the `auto_master` map.**

```
/common-dir      auto_common-dir      -nosuid
```

The `auto_common-dir` map determines the contents of the common directory.

- 3. Add the `-nosuid` option as a precaution to prevent users from running the `setuid` programs that might exist in any workspace.**

- 4. Add entries to the `auto_common-dir` map.**

The `auto_common-dir` map is organized so that each entry describes a subproject. Your first attempt yields a map that resembles the following:

```
project1  system1:/export/common-dir/&
project2  system1:/export/common-dir/&
app1      system2:/export/common-dir/&
```

The ampersand (&) at the end of each entry is an abbreviation for the entry key.

Example 7 Consolidating Project-Related Files Under `/ws`

Assume that you are the administrator of a large software development project. You plan to make all project-related files available under a directory that is called `/ws`. This directory is to be common across all workstations at the site.

Add an entry for the `/ws` directory to the site `auto_master` map.

```
/ws    auto_ws    -nosuid
```

The `auto_ws` map determines the contents of the `/ws` directory. The `-nosuid` prevents users from running `setuid` programs that might exist in any workspaces. Add entries to the `auto_ws` map such that each entry describes a subproject. Your first attempt yields a map that resembles the following:

```
compiler  alpha:/export/ws/&
windows  alpha:/export/ws/&
files    bravo:/export/ws/&
drivers  alpha:/export/ws/&
man      bravo:/export/ws/&
tools    delta:/export/ws/&
```

The ampersand (&) at the end of each entry is an abbreviation for the entry key. For instance, the first entry is equivalent to the following:

```
compiler  alpha:/export/ws/compiler
```

This first attempt provides a map that appears simple, but additional refinements are necessary. The project organizer decides that the documentation in the `man` entry should be provided as a subdirectory under each subproject. Also, each subproject requires subdirectories to describe several versions of the software. You must assign each of these subdirectories to an entire disk partition on the server.

Modify the entries in the map as follows:

```
compiler \
  /vers1.0  alpha:/export/ws/&/vers1.0 \
  /vers2.0  bravo:/export/ws/&/vers2.0 \
  /man      bravo:/export/ws/&/man
windows \
  /vers1.0  alpha:/export/ws/&/vers1.0 \
  /man      bravo:/export/ws/&/man
files \
  /vers1.0  alpha:/export/ws/&/vers1.0 \
  /vers2.0  bravo:/export/ws/&/vers2.0 \
  /vers3.0  bravo:/export/ws/&/vers3.0 \
  /man      bravo:/export/ws/&/man
drivers \
  /vers1.0  alpha:/export/ws/&/vers1.0 \
```

```
    /man      bravo:/export/ws/&/man
tools \
  /          delta:/export/ws/&
```

Although the map now appears to be much larger, the map still contains only the five entries. Each entry is larger because each entry contains multiple mounts. For instance, a reference to `/ws/compiler` requires three mounts for the `vers1.0`, `vers2.0`, and `man` directories.

The backslash at the end of each line indicates that the entry is continued onto the next line. Effectively, the entry is one long line, though line breaks and some indenting have been used to make the entry more readable.

The `tools` directory contains software development tools for all subprojects, so this directory is not subject to the same subdirectory structure. The `tools` directory continues to be a single mount.

This arrangement provides the administrator with much flexibility. Software projects typically consume substantial amounts of disk space. Through the life of the project, you might be required to relocate and expand various disk partitions. If these changes are reflected in the `auto_ws` map, you do not need to notify the users because the directory hierarchy under `/ws` is not changed.

Because the servers `alpha` and `bravo` view the same `auto_ws` map, any users who log in to these systems can find the `/ws` namespace as expected. These users are provided with direct access to local files through loopback mounts instead of NFS mounts.

▼ How to Set Up Different Architectures to Access a Shared Namespace

You need to assemble a shared namespace for local executables, and applications, such as spreadsheet applications and word-processing packages. The clients of this namespace use several different workstation architectures that require different executable formats. Also, some workstations are running different releases of the operating system.

- 1. Create the `auto_local` map.**

For more information about naming services, see [Working With Oracle Solaris 11.3 Directory and Naming Services: DNS and NIS](#).

- 2. Choose a single, site-specific name for the shared namespace.**

This name makes the files and directories that belong to this space easily identifiable. For example, if you choose `/usr/local` as the name, the path `/usr/local/bin` is clearly a part of this namespace.

3. Create an autofs indirect map and choose a pathname to anchor it.

Mount this map at `/usr/local`. Set up the following entry in the NIS `auto_master` map:

```
/usr/local    auto_local    -ro
```

Notice that the `-ro` mount option means that clients cannot write to any files or directories.

4. Export the appropriate directory on the server.**5. Include a `bin` entry in the `auto_local` map.**

Your directory structure resembles the following:

```
bin    aa:/export/local/bin
```

where, `aa` is the name of the server.

6. (Optional) To serve clients of different architectures, change the entry by adding the autofs CPU variable.

```
bin    aa:/export/local/bin/$CPU
```

where, `aa` is the name of the server.

- For SPARC clients – Place executables in `/export/local/bin/sparc`.
- For x86 clients – Place executables in `/export/local/bin/i386`.

▼ How to Support Incompatible Client Operating System Versions

1. Combine the architecture type with a variable that determines the operating system type of the client.

You can combine the autofs `OSREL` variable with the `CPU` variable to form a name that determines both CPU type and OS release.

2. Create the following map entry.

```
bin    aa:/export/local/bin/$CPU$OSREL
```

For clients that are running Version 5.6 of the operating system, export the following file systems:

- For SPARC clients – Export `/export/local/bin/sparc5.6`.

- For x86 clients – Place executables in `/export/local/bin/i3865.6`.

▼ How to Replicate Shared Files Across Several Servers

The best way to share replicated file systems that are read-only is to use failover. For more information about failover, see [“Client-Side Failover” on page 41](#).

1. **Become an administrator.**

For more information, see [“Using Your Assigned Administrative Rights” in *Securing Users and Processes in Oracle Solaris 11.3*](#).

2. **In the `autofs` maps, create a comma-separated list of all replica servers.**

For example:

```
bin    aa,bb,cc,dd:/export/local/bin/$CPU
```

Autofs chooses the nearest server. If a server has several network interfaces, list each interface. Autofs chooses the nearest interface to the client, avoiding unnecessary routing of NFS traffic.

Autofs Security Restrictions

The `nosuid` option prevents users from creating files with the `setuid` or `setgid` bit set.

This entry overrides the entry for `/home` in a generic local `/etc/auto_master` file. For information about the generic local `/etc/auto_master` file, see [“Setting Up a Common View of `/home`” on page 100](#). The override happens because the `+auto_master` reference to the external name service map occurs before the `/home` entry in the file. If the entries in the `auto_home` map include mount options, the `nosuid` option is overwritten. Therefore, either no options should be used in the `auto_home` map or the `nosuid` option must be included with each entry.

Note - Do not mount the home directory disk partitions on or under `/home` on the server.

▼ How to Apply Autofs Security Restrictions

1. **Become an administrator.**

For more information, see [“Using Your Assigned Administrative Rights” in *Securing Users and Processes in Oracle Solaris 11.3*](#).

2. **Create a `-nosuid` entry in the name service `auto_master` file.**

```
/home    auto_home    -nosuid
```

▼ How to Use a Public File Handle With Autofs

1. **Become an administrator.**

For more information, see [“Using Your Assigned Administrative Rights” in *Securing Users and Processes in Oracle Solaris 11.3*](#).

2. **Create the following entry in the `autofs` map.**

```
/usr/local    -ro,public    bee:/export/share/local
```

The `public` option forces the public handle to be used. If the NFS server does not support a public file handle, the mount fails.

▼ How to Use NFS URLs With Autofs

1. **Become an administrator.**

For more information, see [“Using Your Assigned Administrative Rights” in *Securing Users and Processes in Oracle Solaris 11.3*](#).

2. **Create the following entry in the `autofs` map.**

```
/usr/local    -ro    nfs://server-name/export/share/local
```

The service tries to use the public file handle on the NFS server. However, if the server does not support a public file handle, the MOUNT protocol is used.

Disabling Autofs Browsability

The default version of `/etc/auto_master` that is installed has the `-nobrowse` option added to the entries for `/home` and `/net`. In addition, the upgrade procedure adds the `-nobrowse` option to the `/home` and `/net` entries in `/etc/auto_master` if these entries have not been modified.

However, you might have to make these changes manually or to turn off browsability for site-specific autofs mount points after the installation.

This section describes how to turn off the browsability feature for a single client, all clients, and a selected file system.

▼ How to Completely Disable Autofs Browsability on a Single NFS Client

1. **Become an administrator on the NFS client.**

For more information, see [“Using Your Assigned Administrative Rights”](#) in *Securing Users and Processes in Oracle Solaris 11.3*.

2. **Change the autofs SMF configuration parameter.**

```
$ sharectl set -p nobrowse=TRUE autofs
```

3. **Restart the autofs service.**

```
$ svcadm restart system/filesystem/autofs
```

▼ How to Disable Autofs Browsability for All Clients

To disable browsability for all clients, you must employ a name service such as NIS. Otherwise, you need to manually edit the automounter maps on each client. In this example, the browsability of the /home directory is disabled. You must follow this procedure for each indirect autofs node that needs to be disabled.

1. **Add the -nobrowse option to the /home entry in the name service auto_master file.**

```
/home    auto_home    -nobrowse
```

2. **Run the automount command on all clients to make the new behaviour effective.**

The new behavior also becomes effective after a reboot.

```
$ /usr/sbin/automount
```

▼ How to Disable Autofs Browsability on a Selected File System

In this example, browsability of the /net directory is disabled. You can use the same procedure for /home or any other autofs mount points.

1. Verify the search order for the automount naming services.

The `config/automount` property in the `name-service/switch` service shows the search order for the automount information.

```
$ svcprop -p config svc:/system/name-service/switch
config/value_authorization astring solaris.smf.value.name-service.switch
config/printer astring user\ files
config/default astring files\ nis
config/automount astring files\ nis
```

The last entry shows that local automount files are searched first and then the NIS service is checked. The `config/default` entry specifies the search order for all naming information not specifically listed. If you are using the LDAP naming service, the `config/automount` property in the `name-service/switch` service shows `ldap` in the search order for the automount information.

```
$ svcprop -p config svc:/system/name-service/switch
config/value_authorization astring solaris.smf.value.name-service.switch
config/printer astring user\ files\ ldap
config/default astring files\ ldap
config/automount astring files\ ldap
```

2. Check the position of the `+auto_master` entry in `/etc/auto_master`.

For additions to the local files to have precedence over the entries in the namespace, the `+auto_master` entry must be moved to follow `/net`.

```
# Master map for automounter
#
/net    -hosts    -nosuid
/home   auto_home
/nfs4   -fedfs     -ro,nosuid,nobrowse
+auto_master
```

A standard configuration places the `+auto_master` entry at the top of the file, which prevents any local changes from being used.

3. Add the `nobrowse` option to the `/net` entry in the `/etc/auto_master` file.

```
/net    -hosts    -nosuid,nobrowse
```

4. On all clients, run the `automount` command.

The new behavior becomes effective after running the `automount` command on the client systems or after a reboot.

```
$ /usr/sbin/automount
```


Commands for Managing Network File Systems

This chapter describes the command-line utilities that are used to manage network file systems. For a quick reference to commonly used commands for NFS file sharing and troubleshooting, see [Appendix A, “NFS File Sharing Command Reference”](#).

This chapter contains the following topics:

- “NFS Commands” on page 111
- “FedFS Commands” on page 132

Note - If your system has zones enabled and you want to use this feature in a non-global zone, see [Introduction to Oracle Solaris Zones](#).

NFS Commands

These commands must be run as root to be fully effective, but requests for information can be made by all users:

- “[automount Command](#)” on page 112
- “[clear_locks Command](#)” on page 113
- “[fsstat Command](#)” on page 113
- “[mount Command](#)” on page 114
- “[mountall Command](#)” on page 121
- “[nfsref Command](#)” on page 132
- “[sharectl Command](#)” on page 122
- “[share Command](#)” on page 124
- “[shareall Command](#)” on page 130
- “[showmount Command](#)” on page 130
- “[umount Command](#)” on page 120
- “[umountall Command](#)” on page 121

- [“unshare Command” on page 129](#)
- [“unshareall Command” on page 130](#)

In addition, commands associated with the FedFS service are covered in [“FedFS Commands” on page 132](#).

automount Command

This command installs `autofs` mount points and associates the information in the automaster files with each mount point. The syntax of the command is as follows:

```
automount [-t duration] [-v]
```

- | | |
|---------------------------------|--|
| <code>-t <i>duration</i></code> | Sets the time, in seconds, that a file system is to remain mounted. |
| <code>-v</code> | Selects the verbose mode. Running this command in verbose mode enables easier troubleshooting. |

If not specifically set, the value for duration is set to 5 minutes. In most circumstances, this value is good. However, on systems that have many automounted file systems, you might need to increase the duration value. In particular, if a server has many users active, checking the automounted file systems every 5 minutes can be inefficient. Checking the `autofs` file systems every 1800 seconds, which is 30 minutes, could be more optimal. By not unmounting the file systems every 5 minutes, `/etc/mnttab` can become large. To reduce the output when `df` checks each entry in `/etc/mnttab`, you can filter the output from `df` by using the `-F` option (see the [df\(1M\)](#) man page) or by using `egrep`.

You should consider that adjusting the duration also changes how quickly changes to the automounter maps are reflected. Changes cannot be seen until the file system is unmounted. Refer to [“Modifying the Maps” on page 98](#) for instructions on how to modify automounter maps.

You can make the same specifications with the `sharectl` command that you would make on the command line. However, unlike the command-line options, the SMF repository preserves your specifications, through service restarts and system reboots, as well as system upgrades. You can set the following parameters for the `automount` command.

```
timeout
```

Sets the duration for a file system to remain idle before the file system is unmounted. This keyword is the equivalent of the `-t` argument for the `automount` command. The default value is 600.

`automount_verbose`

Provides notification of `autofs` mounts, unmounts, and other nonessential events. This keyword is the equivalent of the `-v` argument for `automount`. The default value is `FALSE`.

clear_locks Command

This command enables you to remove all file, record, and share locks for an NFS client. You must be `root` to run this command. From an NFS server, you can clear the locks for a specific client. From an NFS client, you can clear locks for that client on a specific server. The following example would clear the locks for an NFS client named `tulip` on the current system.

```
$ clear_locks tulip
```

Use the `-s` option to specify which NFS host to clear the locks from. You must run this option from the NFS client that created the locks. In this situation, the locks from the client would be removed from an NFS server named `bee`.

```
$ clear_locks -s bee
```



Caution - This command should be run only when a client crashes and cannot clear its locks. To avoid data corruption problems, do not clear locks for an active client.

fsstat Command

The `fsstat` utility enables you to monitor file system operations by file system type and by mount point. Various options enable you to customize the output:

- `-i` Displays statistics about the I/O operations for mount points
- `-n` Displays statistics about the naming operations for mount points

The following example shows output for NFS Version 3, NFS Version 4, and the root mount point.

```
$ fsstat nfs3 nfs4 /
new      name  name  attr  attr  lookup  rddir  read  read  write  write
file     remov chng  get   set   ops     ops    ops  bytes ops   bytes
3.81K    90   3.65K 5.89M 11.9K 35.5M   26.6K 109K 118M 35.0K 8.16G
nfs3
```

```

759      503      457    93.6K    1.44K    454K    8.82K    65.4K    827M    292    223K
nfs4
25.2K    18.1K    1.12K    54.7M    1017    259M    1.76M    22.4M    20.1G    1.43M    3.77G /

```

The following example uses the `-i` option to provide statistics about the I/O operations for NFS Version 3, NFS Version 4, and the root mount point.

```

$ fsstat -i nfs3 nfs4 /
read    read    write   write   rddir   rddir   rwlock  rwunlock
ops     bytes   ops     bytes   ops     bytes   ops     ops
109K    118M   35.0K   8.16G   26.6K   4.45M   170K    170K   nfs3
65.4K   827M   292     223K   8.82K   2.62M   74.1K   74.1K   nfs4
22.4M   20.1G   1.43M   3.77G   1.76M   3.29G   25.5M   25.5M   /

```

The following example uses the `-n` option to provide statistics about the naming operations for NFS Version 3, NFS Version 4, and the root mount point.

```

$ fsstat -n nfs3 nfs4 /
lookup  creat   remov  link   renam  mkdir  rmdir  rddir  symlnk  rdlnk
35.5M   3.79K   90     2     3.64K  5      0      26.6K   11     136K   nfs3
454K    403    503    0     101    0      0      8.82K   356    1.20K   nfs4
259M    25.2K   18.1K  114   1017   10     2      1.76M   12     8.23M   /

```

For more information, see the [fsstat\(1M\)](#) man page.

mount Command

With this command, you can attach a named file system, either local or remote, to a specified mount point. For more information, see the [mount\(1M\)](#) man page. Used without arguments, `mount` displays a list of file systems that are currently mounted on your computer.

Each file system type included in the standard Oracle Solaris installation has specific options for the `mount` command. For NFS file systems options, see the [mount_nfs\(1M\)](#) man page. For UFS file system options, see the [mount_ufs\(1M\)](#) man page.

You can select a path name to mount from an NFS server by using an NFS URL instead of the standard `server:/pathname` syntax. See [“How to Mount an NFS File System by Using an NFS URL” on page 79](#) for further information.



Caution - The `mount` command does not warn about invalid options. The command silently ignores any options that cannot be interpreted. Ensure that you verify all of the options that were used so that you can prevent unexpected behavior.

mount Options for NFS File Systems

This section describes some of the options that can follow the `-o` flag when you are mounting an NFS file system. For a complete list of options, refer to the [mount_nfs\(1M\)](#) man page.

`bg|fg`

These options can be used to select the retry behavior if a mount fails. The `bg` option causes the mount attempts to be run in the background. The `fg` option causes the mount attempt to be run in the foreground. The default is `fg`, which is the best selection for file systems that must be available because it prevents further processing until the mount is complete. `bg` is a good selection for noncritical file systems because the client can do other processing while waiting for the mount request to be completed.

`forcedirectio`

This option improves performance of large sequential data transfers. Data is copied directly to a user buffer. No caching is performed in the kernel on the client. This option is off by default (`noforcedirectio`).

To permit an application to issue concurrent writes, as well as concurrent reads and writes, to a single file on the client, use the `forcedirectio` mount option. This option, enables this functionality for all files within the mounted file system. You could also enable this functionality on a single file on the client by using the `directio()` interface. Unless this functionality has been enabled, writes to files are serialized. Also, if concurrent writes or concurrent reads and writes are occurring, then POSIX semantics are no longer being supported for that file.

For an example of how to use this option, refer to [“Using the mount Command” on page 117](#).

`largefiles`

With this option, you can access files that are larger than 2 GB. Whether a large file can be accessed can only be controlled on the server, so this option is silently ignored on NFS Version 3 mounts. By default, all UFS file systems are mounted with `largefiles`. For mounts that use the NFS Version 2 protocol, the `largefiles` option causes the mount to fail with an error.

`nolargefiles`

This option for UFS mounts guarantees that no large files can exist on the file system. Because the existence of large files can be controlled only on the NFS server, no option for `nolargefiles` exists when using NFS mounts. Attempts to NFS-mount a file system by using this option are rejected with an error.

nosuid|suid

The **nosuid** option is the equivalent of specifying the **nodevices** option with the **nosetuid** option. When the **nodevices** option is specified, the opening of device-special files on the mounted file system is disallowed. When the **nosetuid** option is specified, the **setuid** bit and **setgid** bit in binary files that are located in the file system are ignored. The processes run with the privileges of the user who executes the binary file.

The **suid** option is the equivalent of specifying the **devices** option with the **setuid** option. When the **devices** option is specified, the opening of device-special files on the mounted file system is allowed. When the **setuid** option is specified, the **setuid** bit and the **setgid** bit in binary files that are located in the file system are honored by the kernel.

If neither option is specified, the default option is **suid**, which provides the default behavior of specifying the **devices** option with the **setuid** option.

The following table describes the effect of combining **nosuid** or **suid** with **devices** or **nodevices**, and **setuid** or **nosetuid**. Note that in each combination of options, the most restrictive option determines the behavior.

| Behavior From the Combined Options | Option | Option | Option |
|---|---------------|-----------------|------------------|
| The equivalent of nosetuid with nodevices | nosuid | nosetuid | nodevices |
| The equivalent of nosetuid with devices | nosuid | nosetuid | devices |
| The equivalent of nosetuid with nodevices | nosuid | setuid | nodevices |
| The equivalent of nosetuid with devices | nosuid | setuid | devices |
| The equivalent of nosetuid with nodevices | suid | nosetuid | nodevices |
| The equivalent of nosetuid with devices | suid | nosetuid | devices |
| The equivalent of setuid with nodevices | suid | setuid | nodevices |
| The equivalent of setuid with devices | suid | setuid | devices |

The **nosuid** option provides additional security for NFS clients that access potentially untrusted servers. Mounting remote file systems with this option reduces the chance of privilege escalation through importing untrusted devices or importing untrusted **setuid** binary files. All these options are available in all Oracle Solaris file systems.

public

This option forces the use of the **public** file handle when contacting the NFS server. If the **public** file handle is supported by the server, the mounting operation is faster because the

MOUNT protocol is not used. Also, because the MOUNT protocol is not used, the public option allows mounting to occur through a firewall.

`rw|ro`

The `-rw` and `-ro` options indicate whether a file system is to be mounted read-write or read-only. The default is read-write, which is the appropriate option for remote home directories, mail-spooling directories, or other file systems that need to be changed by users. The read-only option is appropriate for directories that should not be changed by users. For example, shared copies of the man pages should not be writable by users.

`sec=mode`

You can use this option to specify the authentication mechanism to be used during the mount transaction. The available values for *mode* are:

- `krb5` for Kerberos Version 5 authentication service
- `krb5i` for Kerberos Version 5 with integrity
- `krb5p` for Kerberos Version 5 with privacy
- `none` for no authentication
- `dh` for Diffie-Hellman (DH) authentication
- `sys` for standard UNIX authentication

The modes are also defined in `/etc/nfssec.conf`.

`soft|hard`

An NFS file system that is mounted with the `soft` option returns an error if the server does not respond. The `hard` option causes the mount to continue to retry until the server responds. The default is `hard`, which should be used for most file systems. Applications frequently do not check return values from `soft`-mounted file systems, which can make the application fail or can lead to corrupted files. If the application does check the return values, routing problems and other conditions can still confuse the application or lead to file corruption. In most situations, the `soft` option should not be used. If a file system is mounted by using the `hard` option and becomes unavailable, an application that uses this file system hangs until the file system becomes available.

Using the `mount` Command

The following examples show different scenarios:

- In NFS Version 2 or NFS Version 3, both of the following commands mount an NFS file system from the server `bee` read-only.

```
$ mount -F nfs -r bee:/export/share/man /usr/man
```

```
$ mount -F nfs -o ro bee:/export/share/man /usr/man
```

In NFS Version 4, the following command line would accomplish the same mount.

```
$ mount -F nfs -o vers=4 -r bee:/export/share/man /usr/man
```

- In NFS Version 2 or NFS Version 3, the `-O` option in the following command forces the man pages from the server bee to be mounted on the local system even if `/usr/man` has already been mounted.

```
$ mount -F nfs -O bee:/export/share/man /usr/man
```

In NFS Version 4, the following command would accomplish the same mount:

```
$ mount -F nfs -o vers=4 -O bee:/export/share/man /usr/man
```

- In NFS Version 2 or NFS Version 3, the following command uses client failover.

```
$ mount -F nfs -r bee,wasp:/export/share/man /usr/man
```

In NFS Version 4, the following command uses client failover.

```
$ mount -F nfs -o vers=4 -r bee,wasp:/export/share/man /usr/man
```

Note - When used from the command line, the listed servers must support the same version of the NFS protocol. Do not use both NFS Version 2 and NFS Version 3 servers when running `mount` from the command line. You can use both servers with `auto fs` because `auto fs` automatically selects the best subset of NFS Version 2 or NFS Version 3 servers.

- The following example shows how to use an NFS URL with the `mount` command in NFS Version 2 or NFS Version 3.

```
$ mount -F nfs nfs://bee:/export/share/man /usr/man
```

The following example shows how to use an NFS URL with the `mount` command in NFS Version 4.

```
$ mount -F nfs -o vers=4 nfs://bee:/export/share/man /usr/man
```

- The following example shows how to use the `quota` mount option to enable the user to check the file system disk quota and usage.

```
$ mount -F nfs -o quota bee:/export/share/man /usr/man
```

- The following example shows how to use the `forcedirectio` mount option to enable the client to permit concurrent writes, as well as concurrent reads and writes, to a file.

```
$ mount -F nfs -o forcedirectio bee:/home/somebody /mnt
```

In this example, the command mounts an NFS file system from the server `bee` and enables concurrent reads and writes for each file in the directory `/mnt`. When support for concurrent reads and writes is enabled, the following occurs.

- The client permits applications to write to a file in parallel.
- Caching is disabled on the client. Consequently, data from reads and writes is kept on the server. More explicitly, because the client does not cache the data that is read or written, any data that the application does not already have cached for itself is read from the server. The client's operating system does not have a copy of this data. Normally, the NFS client caches data in the kernel for applications to use.

Because caching is disabled on the client, the read-ahead and write-behind processes are disabled. A read-ahead process occurs when the kernel anticipates the data that an application might request next. The kernel then starts the process of gathering that data in advance. The kernel's goal is to have the data ready before the application makes a request for the data.

The client uses the write-behind process to increase write throughput. Instead of immediately starting an I/O operation every time an application writes data to a file, the data is cached in memory. Later, the data is written to the disk.

Potentially, the write-behind process permits the data to be written in larger chunks or to be written asynchronously from the application. Typically, the result of using larger chunks is increased throughput. Asynchronous writes permit overlap between application processing and I/O processing. Also, asynchronous writes permit the storage subsystem to optimize the I/O by providing a better sequencing of the I/O. Synchronous writes force a sequence of I/O on the storage subsystem that might not be optimal.

- Significant performance degradation can occur if the application is not prepared to handle the semantics of data that is not being cached. Multithreaded applications avoid this problem.

Note - If support for concurrent writes is not enabled, all write requests are serialized. When a write request is in progress, a second write request has to wait for the first write request to be completed before proceeding.

- The following example shows how to use the `mount` command with no arguments to display file systems that are mounted on a client.

```
$ mount
/ on /dev/dsk/c0t3d0s0 read/write/setuid on Wed Apr 7 13:20:47 2004
/usr on /dev/dsk/c0t3d0s6 read/write/setuid on Wed Apr 7 13:20:47 20041995
/proc on /proc read/write/setuid on Wed Apr 7 13:20:47 2004
/dev/fd on fd read/write/setuid on Wed Apr 7 13:20:47 2004
/tmp on swap read/write on Wed Apr 7 13:20:51 2004
```

```
/opt on /dev/dsk/c0t3d0s5 setuid/read/write on Wed Apr 7 13:20:51 20041995
/home/kathys on bee:/export/home/bee7/kathys
intr/noquota/nosuid/remote on Wed Apr 24 13:22:13 2004
```

umount Command

The `umount` command enables you to remove a remote file system that is currently mounted. You can use the following options with the `umount` command:

| | |
|-----------------|--|
| <code>-v</code> | Enables testing |
| <code>-a</code> | Unmounts several file systems at one time. If <i>mount-points</i> are included with the <code>-a</code> option, those file systems are unmounted. If no mount points are included, an attempt is made to unmount all file systems that are listed in <code>/etc/mnttab</code> except for the "required" file systems, such as <code>/</code> , <code>/usr</code> , <code>/var</code> , <code>/proc</code> , <code>/dev/fd</code> , and <code>/tmp</code> . Because the file system is already mounted and should have an entry in <code>/etc/mnttab</code> , you do not need to include a flag for the file system type. |
| <code>-f</code> | Forces a busy file system to be unmounted. You can use this option to unhang a client that is hung while trying to mount an unmountable file system. |



Caution - By forcing an unmount of a file system, you can cause data loss if files are being written to that system.

EXAMPLE 8 Unmounting a File System

The following example unmounts a file system that is mounted on `/usr/man`:

```
$ umount /usr/man
```

EXAMPLE 9 Using Options with `umount`

The following example displays the results of running `umount -a -V`:

```
$ umount -a -V
umount /home/kathys
umount /opt
```



```
umount /home  
umount /net
```

Note that this command does not actually unmount the file systems.

mountall Command

Use the `mountall` command to mount all file systems or a specific group of file systems that are listed in a file system table. The command provides the following options:

- F *FSType* Selects the file system type to be accessed
- r Selects all the remote file systems that are listed in a file system table
- l Selects all the local file systems

Because all file systems that are labeled as NFS file system type are remote file systems, some of these options are redundant. For more information, see the [mountall\(1M\)](#) man page.

The following two examples of user input are equivalent:

```
$ mountall -F nfs  
$ mountall -F nfs -r
```

umountall Command

Use the `umountall` command to unmount a group of file systems. You can use the following options with the `umountall` command:

- k Runs the `fuser -k mount-point` command to kill any processes that are associated with the *mount-point*
- s Indicates that unmount is not to be performed in parallel
- l Specifies that only local file systems are to be used
- r Specifies that only remote file systems are to be used

`-h host` Specifies that all file systems from the named host should be unmounted. You cannot combine the `-h` option with `-l` or `-r`.

The following example unmounts all file systems that are mounted from remote hosts:

```
$ umountall -r
```

The following example unmounts all file systems that are currently mounted from the server `bee`:

```
$ umountall -h bee
```

sharectl Command

This release includes the `sharectl` utility, which is an administrative tool that enables you to configure and manage file-sharing protocols such as NFS. You can use this command to do the following:

- Set client and server operational properties
- Display property values for a specific protocol
- Obtain the status of a protocol

The `sharectl` utility uses the following syntax:

```
$ sharectl subcommand [option] [protocol]
```

The `sharectl` utility supports the following subcommands:

| | |
|---------------------|--|
| <code>set</code> | Defines the properties for a file-sharing protocol. For a list of properties and property values, see the parameters described in the nfs(4) man page. |
| <code>get</code> | Displays the properties and property values for the specified protocol. |
| <code>status</code> | Displays whether the specified protocol is enabled or disabled. If no protocol is specified, the status of all file-sharing protocols is displayed. |

For more information about the `sharectl` utility, see the following:

- [sharectl\(1M\)](#) man page
- [“set Subcommand” on page 123](#)
- [“get Subcommand” on page 123](#)
- [“status Subcommand” on page 124](#)

set Subcommand

The set subcommand, which defines the properties for a file-sharing protocol, supports the following options:

- h Provides an online-help description
- p Defines a property for the protocol

The set subcommand uses the following syntax:

```
$ sharectl set [-h] [-p property=value] protocol
```

You must have root privileges to use the set subcommand.

You do not need to repeat this command for each additional property value. You can use the -p option multiple times to define multiple properties in the same command.

The following example sets the minimum version of the NFS protocol for the client to 3:

```
$ sharectl set -p client_versmin=3 nfs
```

The following example shows how to require clients to use reserved ports for NFS calls for all file systems shared with AUTH_SYS:

```
$ sharectl set -p resvport=on nfs
```

get Subcommand

The get subcommand, which displays the properties and property values for the specified protocol, supports the following options:

- h Provides an online-help description.
- p Identifies the property value for the specified property. If the -p option is not used, all property values are displayed.

The get subcommand uses the following syntax:

```
$ sharectl get [-h] [-p property] protocol
```

You must have root privileges to use the get subcommand.

The following example uses servers, which is the property that enables you to specify the maximum number of concurrent NFS requests:

```
$ sharectl get -p servers nfs
servers=1024
```

In the following example, because the `-p` option is not used, all property values are displayed:

```
$ sharectl get nfs
servers=1024
listen_backlog=32
protocol=ALL
servers=32
lockd_listen_backlog=32
lockd_servers=20
lockd_retransmit_timeout=5
grace_period=90
nfsmapid_domain=example.com
server_versmin=2
server_versmax=4
client_versmin=2
client_versmax=4
server_delegation=on
max_connections=-1
device=
```

status Subcommand

The `status` subcommand displays whether the specified protocol is enabled or disabled. It supports the `-h` option, which provides an online-help description.

The `status` subcommand uses the following syntax:

```
$ sharectl status [-h] [protocol]
```

The following example shows the status of the NFS protocol:

```
$ sharectl status nfs
nfs    enabled
```

share Command

Use the `share` command to make a local file system on an NFS server available for mounting. You can also use the `share` command to display a list of the file systems on your system that are currently shared. The NFS server must be running for the `share` command to work.

The objects that can be shared include any directory tree. However, each file system hierarchy is limited by the disk slice or partition that the file system is located on.

A file system cannot be shared if that file system is part of a larger file system that is already being shared. For example, if `/usr` and `/usr/local` are on one disk slice, `/usr` can be shared or `/usr/local` can be shared. However, if both directories need to be shared with different share options, `/usr/local` must be moved to a separate disk slice.

You can gain access to a file system that is read-only shared through the file handle of a file system that is read-write shared. However, the two file systems have to be on the same disk slice. To create a more secure situation, place those file systems that need to be read-write on a separate partition or separate disk slice from the file systems that you need to share as read-only.

Note - For information about how NFS Version 4 functions when a file system is unshared and then reshared, refer to [“Unsharing and Resharing a File System in NFS Version 4” on page 29](#).

share Options

Some of the options that you can include with the `-o` flag are as follows:

`rw|ro`

The *pathname* file system is shared read-write or read-only for all clients.

`rw=access-list`

The file system is shared read-write for the clients that are listed only. All other requests are denied. See [“Setting Access Lists With the share Command” on page 128](#) for more information. You can use this option to override an `-ro` option.

NFS-Specific share Options

The options that you can use with NFS file systems include the following:

`aclok`

This option enables an NFS server that supports the NFS Version 2 protocol to be configured to do access control for NFS Version 2 clients. Without this option, all clients are given minimal access. With this option, the clients have maximal access. For instance, on file systems that are shared with the `aclok` option, if anyone has read permissions, everyone does. However, without this option, you can deny access to a client who should have access permissions. A decision to permit too much access or too little access depends

on the security systems already in place. See [“Using Access Control Lists to Protect UFS Files”](#) in *Securing Files and Verifying File Integrity in Oracle Solaris 11.3* for more information about access control lists (ACLs).

Note - To use ACLs, ensure that clients and servers run software that supports the NFS Version 3 and NFS_ACL protocols. If the software only supports the NFS Version 3 protocol, clients obtain correct access but cannot manipulate the ACLs. If the software supports the NFS_ACL protocol, the clients obtain correct access and can manipulate the ACLs.

`anon=uid`

You use `anon` to select the user ID of unauthenticated users. If you set `anon` to `-1`, the server denies access to unauthenticated users. Because granting root access by setting `anon=0` allows unauthenticated users to have root access, use the `root` option instead.

`index=filename`

When a user accesses an NFS URL, the `index=filename` option forces the HTML file to load instead of displaying a list of the directory. This option mimics the action of current browsers if an `index.html` file is found in the directory that the HTTP URL is accessing. This option is the equivalent of setting the `DirectoryIndex` option for `httpd`. For instance, suppose that `share` command reports the following:

```
export_web /export/web  nfs sec=sys,public,index=index.html,ro
```

These URLs then display the same information:

```
nfs://server/dir
nfs://server/dir/index.html
nfs://server//export/web/dir
nfs://server//export/web/dir/index.html
http://server/dir
http://server/dir/index.html
```

`log=tag`

This option specifies the tag in `/etc/nfs/nfslog.conf` that contains the NFS server logging configuration information for a file system. This option must be selected to enable NFS server logging.

`nosuid`

This option signals that all attempts to enable the `setuid` or `setgid` mode should be ignored. NFS clients cannot create files with the `setuid` or `setgid` bits on.

`public`

The `public` option has been added to the `share` command to enable WebNFS browsing. Only one file system on a server can be shared with this option.

`root=access-list`

The server gives root access to the hosts in the list. By default, the server does not give root access to any remote hosts. If the selected security mode is anything other than `sec=sys`, you can only include client host names in the list. See [“Setting Access Lists With the share Command” on page 128](#) for more information.



Caution - Granting root access to other hosts has wide security implications. Use the `-root=` option with extreme caution.

`root=client-name`

The *client-name* value is used with `AUTH_SYS` authentication to check the client's IP address against a list of addresses provided by `exportfs(1B)`. If a match is found, root access is given to the file systems being shared.

`root=hostname`

For secure NFS modes such as `AUTH_SYS` or `RPCSEC_GSS`, the server checks the clients' principal names against a list of host-based principal names that are derived from an access list. The generic syntax for the client's principal name is `root@hostname`. For Kerberos V, the syntax is `root/hostname.fully.qualified@REALM`. When you use the *hostname* value, the clients on the access list must have the credentials for a principal name. For Kerberos V, the client must have a valid keytab entry for its `root/hostname.fully.qualified@REALM` principal name. For more information, see [“Configuring Kerberos Clients” in *Managing Kerberos and Other Authentication Services in Oracle Solaris 11.3*](#).

`sec=mode[:mode]`

This option sets the security modes that are needed to obtain access to the file system. By default, the security mode is UNIX authentication. You can specify multiple modes, but use each security mode only once per command line. Each `sec=` option applies to any subsequent `rw`, `ro`, `rw=`, `ro=`, `root=`, and `window=` options until another `sec=` is encountered. The use of `-sec=none` maps all users to user `nobody`.

`window=value`

value selects the maximum lifetime in seconds of a credential on the NFS server. The default value is 30000 seconds or 8.3 hours.

`resvport`

This option enforces the use of reserved port for individual file systems.

The following example shows how to require the client to use reserved port for an NFS share that is shared with `AUTH_SYS`:

```
$ share -F NFS -o resvport=on /usr/src
```

Setting Access Lists With the share Command

The access list that you provide with the `share` command can include a domain name, a subnet number, or an entry to deny access, as well as the standard `-ro=`, `-rw=`, or `-root=` options. These extensions should simplify file access control on a single server without having to change the namespace or maintain long lists of clients.

The following example provides read-only access for most systems but allows read-write access for `rose` and `lilac`:

```
$ share -F nfs -o ro,rw=rose:lilac /usr/src
```

The following example assigns read-only access to any host in the `eng` netgroup. The client `rose` is specifically given read-write access.

```
$ share -F nfs -o ro=eng,rw=rose /usr/src
```

Note - You cannot specify both `rw` and `ro` without arguments. If no read-write option is specified, the default is read-write for all clients.

To share one file system with multiple clients, you must type all options on the same line. If you issue multiple invocations of the `share` command on the same object, only the last command that is run is applied. The following example enables read-write access to three client systems, but only `rose` and `tulip` are given access to the file system as `root`.

```
$ share -F nfs -o rw=rose:lilac:tulip,root=rose:tulip /usr/src
```

When sharing a file system that uses multiple authentication mechanisms, ensure that you include the `-ro`, `-ro=`, `-rw`, `-rw=`, `-root`, and `-window` options after the correct security modes. In this example, UNIX authentication is selected for all hosts in the netgroup that is named `eng`. These hosts can mount the file system only in read-only mode. The hosts `tulip` and `lilac` can mount the file system read-write if these hosts use Diffie-Hellman authentication. With these options, `tulip` and `lilac` can mount the file system read-only even if these hosts are not using DH authentication. However, the host names must be listed in the `eng` netgroup.

```
$ share -F nfs -o sec=dh,rw=tulip:lilac,sec=sys,ro=eng /usr/src
```

Even though UNIX authentication is the default security mode, UNIX authentication is not included if the `-sec` option is used. Therefore, you must include a `-sec=sys` option if UNIX authentication is to be used with any other authentication mechanism.

You can use a DNS domain name in the access list by preceding the actual domain name with a dot. The string that follows the dot is a domain name, not a fully qualified host name. The following example allows mount access to all hosts in the `eng.example.com` domain:


```
$ share -F nfs -o ro=:.eng.example.com /export/share/man
```

In this example, the single dot matches all hosts that are matched through the NIS namespace. The results that are returned from these name services do not include the domain name. The `.eng.example.com` entry matches all hosts that use DNS for namespace resolution. Because DNS always returns a fully qualified host name, the longer entry is required if you use a combination of DNS and the other namespaces.

You can use a subnet number in an access list by preceding the actual network number or the network name with an at (@) sign. This character differentiates the network name from a netgroup or a fully qualified host name. You must identify the subnet in either `/etc/networks` or in an NIS namespace. The following entries have the same effect if the `192.0.2` subnet has been identified as the `eng` network:

```
$ share -F nfs -o ro=@eng /export/share/man
$ share -F nfs -o ro=@192.0.2 /export/share/man
$ share -F nfs -o ro=@192.0.2.0 /export/share/man
```

The last two entries show that you do not need to include the full network address.

If the network prefix is not byte aligned, as with Classless Inter-Domain Routing (CIDR), the mask length can be explicitly specified on the command line. The mask length is defined by following either the network name or the network number with a slash and the number of significant bits in the prefix of the address. For example:

```
$ share -f nfs -o ro=@eng/23 /export/share/man
$ share -F nfs -o ro=@192.0.2/23 /export/share/man
```

In these examples, the `/23` indicates that the first 23 bits in the address are to be used as the mask. For additional information about CIDR, see RFC 1519.

You can also select negative access by placing a `-` before the entry. Note that the entries are read from left to right. Therefore, you must place the negative access entries before the entry that the negative access entries apply to:

```
$ share -F nfs -o ro=-rose:.eng.example.com /export/share/man
```

This example would allow access to any hosts in the `eng.example.com` domain except the host that is named `rose`.

unshare Command

The `unshare` command enables you to make a previously available file system unavailable for mounting by clients. When you unshare an NFS file system, access from clients with existing

mounts is inhibited. The file system might still be mounted on the client but the files are not accessible. The `unshare` command deletes the share permanently unless the `-t` option is used to temporarily unshare the file system.

Note - For information about how NFS Version 4 functions when a file system is unshared and then reshared, refer to [“Unsharing and Resharing a File System in NFS Version 4” on page 29](#).

The following example unshares the file system `/usr/src`:

```
$ unshare /usr/src
```

shareall Command

The `shareall` command enables the sharing of multiple file systems. When used with no options, the command shares all entries in the SMF repository. You can include a file name to specify the name of a file that lists share command lines.

The following example shares all file systems that are listed in a local file:

```
$ shareall /etc/dfs/special_dfstab
```

unshareall Command

The `unshareall` command makes all currently shared resources unavailable. The `-F FSType` option selects a list of file system types that are defined in `/etc/dfs/fstypes`. This flag enables you to choose only certain types of file systems to be unshared. The default file system type is defined in `/etc/dfs/fstypes`. To choose specific file systems, use the `unshare` command.

The following example unshares all NFS-type file systems:

```
$ unshareall -F nfs
```

showmount Command

Use the `showmount` command to display the following information:

- All clients that have remotely mounted file systems that are shared from an NFS server
- Only the file systems that are mounted by clients

- Shared file systems with client access information

Note - The `showmount` command shows only NFS Version 2 and Version 3 exports. This command does not show NFS Version 4 exports.

The command syntax is as follows:

```
showmount [-ade] [hostname]
```

- a Prints a list of all the remote mounts. Each entry includes the client name and the directory.
- d Prints a list of the directories that are remotely mounted by clients.
- e Prints a list of the files that are shared or are exported.
- hostname* Selects the NFS server to gather the information from.

If *hostname* is not specified, the local host is queried.

The following example lists all clients and the local directories that the clients have mounted:

```
$ showmount -a bee
lilac:/export/share/man
lilac:/usr/src
rose:/usr/src
tulip:/export/share/man
```

The following example lists the directories that have been mounted:

```
$ showmount -d bee
/export/share/man
/usr/src
```

The following example lists file systems that have been shared:

```
$ showmount -e bee
/usr/src            (everyone)
/export/share/man    eng
```

The `nfs_props/showmount_info` property of the `/network/nfs/server:default` service controls how much information is displayed to a client by the `showmount` command. The default value is `full`. If this value is set to `none` then the client will see only those remote file systems on the server that the client can mount. No information about other clients is displayed. See [Example 5, “Restricting File System Information Displayed to Clients,” on page 80](#) for an example of how to change this property.

nfsref Command

Use the `nfsref` command to add, delete, or list NFSv4 referrals. The command syntax is as follows:

```
nfsref add path location [location...]  
nfsref remove path  
nfsref lookup path
```

| | |
|-----------------|--|
| <i>path</i> | Determines the name for the reparse point. |
| <i>location</i> | Identifies one or more NFS or SMB shared file systems to be associated with the reparse point. |

FedFS Commands

The following commands are associated with the FedFS service:

| | |
|-------------------------------|---|
| <code>nsdb-list</code> | Lists all FedFS data stored in the LDAP server. |
| <code>nsdb-nces</code> | Lists the naming contexts on the LDAP server and the relative distinguished name. |
| <code>nsdb-resolve-fsn</code> | Shows the fileset location for the selected fileset name. |
| <code>nsdb-update-nci</code> | Manages distinguished names for FedFS data. |
| <code>nsdbparams</code> | Manages FedFS connections. |

For examples of how these commands are used, see [“Administering FedFS” on page 91](#).

Troubleshooting Network File Systems

This chapter provides information about NFS troubleshooting strategies, procedures, and commands. This chapter also includes information about troubleshooting `autofs` and a list of the NFS error messages with their meanings.

This chapter contains the following topics:

- [“Strategies for NFS Troubleshooting” on page 133](#)
- [“Commands for Troubleshooting NFS Problems” on page 134](#)
- [“NFS Troubleshooting Procedures” on page 140](#)
- [“Troubleshooting Autofs” on page 146](#)
- [“NFS Error Messages” on page 150](#)

Strategies for NFS Troubleshooting

When tracking an NFS problem, remember the main points of possible failure: the NFS server, the NFS client, and the network. Try to isolate each component to find the one that is not working. The `mountd` and `nfsd` daemons must always be running on the server for remote mounts to succeed.

The `-intr` option is set by default for all mounts. If a program hangs with a server not responding message, you can terminate the program with the keyboard interrupt Control-C.

When the network or server has problems, programs that access hard-mounted remote files fail differently than programs that access soft-mounted remote files. Hard-mounted remote file systems cause the client's kernel to retry the requests until the server responds again. Soft-mounted remote file systems cause the client's system calls to return an error after several attempts. Avoid soft mounting because the errors can result in unexpected application errors and data corruption.

When a file system is hard mounted, a program that tries to access the file system hangs if the server fails to respond. In this situation, the NFS system displays the following message on the console:

NFS server *hostname* not responding still trying

When the server finally responds, the following message appears on the console:

NFS server *hostname* ok

A program that accesses a soft-mounted file system whose server is not responding generates the following message:

NFS operation failed for server *hostname*: error # (*error-message*)



Caution - Because of possible errors, do not soft-mount file systems with read-write data or file systems from which executables are run. Writable data could be corrupted if the application ignores the errors. Mounted executables might not load properly and can fail.

Commands for Troubleshooting NFS Problems

This section describes the commands you can use for troubleshooting NFS problems.

nfsstat Command

This command displays statistical information about NFS and RPC connections. Use the following syntax to display NFS server and client statistics:

```
nfsstat [-cmnr sz]
```

- | | |
|----|---|
| -c | Displays client-side information. |
| -m | Displays statistics for each NFS-mounted file system. |
| -n | Displays the NFS information on both the client side and the server side. |
| -r | Displays RPC statistics. |
| -s | Displays the server-side information. |
| -z | Specifies that the statistics should be set to zero. |

If no options are supplied, the `-cnrs` options are used.

Gathering server-side statistics can be important for debugging problems when new software or new hardware is added to the computing environment. Running this command a minimum of once a week, and storing the numbers, provides a good history of previous performance.

EXAMPLE 10 Displaying NFS Server Statistics

```
$ nfsstat -s

Server rpc:
Connection oriented:
calls      badcalls  nullrecv  badlen    xdrCALL   dupchecks dupreqs
719949194  0         0         0         0         58478624  33
Connectionless:
calls      badcalls  nullrecv  badlen    xdrCALL   dupchecks dupreqs
73753609   0         0         0         0         987278    7254

Server NFSv2:
calls      badcalls  referrals referlinks
25733     0         0         0

Server NFSv3:
calls      badcalls  referrals referlinks
132880073 0         0         0

Server NFSv4:
calls      badcalls  referrals referlinks
488884996 4         0         0
Version 2: (746607 calls)
null      getattr  setattr  root      lookup   readlink  read
883 0%    60 0%    45 0%    0 0%     177446 23% 1489 0%  537366 71%
wrcache  write    create   remove    rename   link      symlink
0 0%     1105 0%   47 0%    59 0%    28 0%    10 0%    9 0%
mkdir    rmdir   readdir  statfs
26 0%    0 0%     27926 3%  108 0%

Version 3: (728863853 calls)
null      getattr  setattr  lookup   access
1365467 0%    496667075 68% 8864191 1%  66510206 9%  19131659 2%
readlink  read     write    create   mkdir
414705 0%    80123469 10% 18740690 2%  4135195 0%  327059 0%
symlink   mknod   remove   rmdir    rename
101415 0%    9605 0%   6533288 0%  111810 0%  366267 0%
link      readdir readdirplus fsstat  fsinfo
2572965 0%   519346 0%   2726631 0%  13320640 1%  60161 0%
pathconf  commit
13181 0%   6248828 0%
```

```

Version 4: (54871870 calls)
null          compound
266963 0%     54604907 99%
Version 4: (167573814 operations)
reserved      access          close          commit
0 0%          2663957 1%     2692328 1%   1166001 0%
create        delegpurge      delegreturn   getattr
167423 0%     0 0%          1802019 1%   26405254 15%
getfh         link           lock          lockt
11534581 6%   113212 0%     207723 0%    265 0%
locku         lookup         lookupp       nverify
230430 0%     11059722 6%   423514 0%    21386866 12%
open          openattr       open_confirm  open_downgrade
2835459 1%    4138 0%      18959 0%     3106 0%
putfh         putpubfh      putrootfh     read
52606920 31%  0 0%         35776 0%    4325432 2%
readdir       readlink      remove        rename
606651 0%     38043 0%     560797 0%    248990 0%
renew         restorefh     savefh        secinfo
2330092 1%    8711358 5%   11639329 6%  19384 0%
setattr       setclientid   setclientid_confirm verify
453126 0%     16349 0%     16356 0%    2484 0%
write         release_lockowner illegal
3247770 1%    0 0%         0 0%
Server nfs_acl:
Version 2: (694979 calls)
null      getacl      setacl      getattr      access      getxattrdir
0 0%      42358 6%    0 0%        584553 84%  68068 9%    0 0%
Version 3: (2465011 calls)
null      getacl      setacl      getxattrdir
0 0%      1293312 52%  1131 0%     1170568 47%

```

The example shows how to display the statistics for RPC and NFS activities. In both sets of statistics, knowing the average number of `badcalls` or `calls` and the number of calls per week can help identify a problem. The `badcalls` value reports the number of bad messages from a client. This value can indicate network hardware problems.

Some of the connections generate write activity on the disks. A sudden increase in these statistics could indicate trouble and should be investigated. For NFS Version 2 statistics, the connections to note are `setattr`, `write`, `create`, `remove`, `rename`, `link`, `symlink`, `mkdir`, and `rmdir`. For NFS Version 3 and NFS Version 4 statistics, the value to watch is `commit`. If the `commit` level is high in one NFS server compared to another almost identical server, check that the NFS clients have enough memory. The number of `commit` operations on the server grows when clients do not have available resources.

pstack Command

The `pstack` command displays a stack trace for each process. The `pstack` command must be run by the owner of the process or by root. You can use the `pstack` command to determine where a process is hung. The only option that is allowed with this command is the process ID of the process that you want to check. For more information about the `pstack` command, see the [proc\(1\)](#) man page.

EXAMPLE 11 Displaying Stack Trace for NFS Process

```
$ /usr/bin/pgrep nfsd
243
$ /usr/bin/pstack 243
243: /usr/lib/nfs/nfsd -a 16
ef675c04 poll (24d50, 2, ffffffff)
000115dc ???????? (24000, 132c4, 276d8, 1329c, 276d8, 0)
00011390 main (3, effffff14, 0, 0, ffffffff, 400) + 3c8
00010fb0 _start (0, 0, 0, 0, 0, 0) + 5c
```

The example shows that the process is waiting for a new connection request, which is a normal response. If the stack shows that the process is still in `poll` after a request is made, the process might be hung. For more information about fixing a hung process, see [“How to Restart NFS Service” on page 145](#). For more information about troubleshooting NFS, see [“NFS Troubleshooting Procedures” on page 140](#).

rpcinfo Command

The `rpcinfo` command generates information about the RPC service that is running on a system. Use the following command syntaxes to display information about the RPC service:

```
rpcinfo [-m|-s] [hostname]
rpcinfo [-T transport] [hostname] [progname]
rpcinfo [-t|-u] [hostname] [progname]
```

| | |
|---------------------------|---|
| <code>-m</code> | Displays a table of statistics of the <code>rpcbind</code> operations |
| <code>-s</code> | Displays a concise list of all registered RPC programs |
| <code>-T transport</code> | Displays information about services that use specific transports or protocols |

| | |
|------------------------|--|
| <code>-t</code> | Probes the RPC programs that use TCP |
| <code>-u</code> | Probes the RPC programs that use UDP |
| <code>transport</code> | Specifies the transport or protocol for the services |
| <code>hostname</code> | Specifies the host name of the server |
| <code>progname</code> | Specifies the name of the RPC program |

For more information about the available options, see the [rpcinfo\(1M\)](#) man page.

If no value is given for `hostname`, the local host name is used. You can substitute the RPC program number for `progname`, but the name is more commonly used. You can use the `-p` option in place of the `-s` option on those systems that do not run the NFS Version 3 software.

The data that is generated by this command can include the following:

- RPC program number
- Version number for a specific program
- Transport protocol in use
- Name of the RPC service
- Owner of the RPC service

EXAMPLE 12 Displaying RPC Service Information

```
$ rpcinfo -s bee |sort -n
  program version(s) netid(s)          service  owner
  100000  2,3,4    udp6,tcp6,udp,tcp,ticlts,ticotsord,ticots portmapper superuser
  100001  4,3,2    udp6,udp,ticlts                          rstatd   superuser
  100003  4,3,2    tcp,udp,tcp6,udp6                        nfs      1
  100005  3,2,1    ticots,ticotsord,tcp,tcp6,ticlts,udp,udp6 mountd   superuser
  100007  1,2,3    ticots,ticotsord,ticlts,tcp,udp,tcp6,udp6 ypbind  1
  100011  1        udp6,udp,ticlts                          rquotad  superuser
  100021  4,3,2,1  tcp,udp,tcp6,udp6                        nlockmgr 1
  100024  1        ticots,ticotsord,ticlts,tcp,udp,tcp6,udp6 status   superuser
  100068  5,4,3,2  ticlts                                    -        superuser
  100083  1        ticotsord                                  -        superuser
  100133  1        ticots,ticotsord,ticlts,tcp,udp,tcp6,udp6 -        superuser
  100134  1        ticotsord                                  -        superuser
  100155  1        ticotsord                                  smserverd superuser
  100169  1        ticots,ticotsord,ticlts                   -        superuser
  100227  3,2      tcp,udp,tcp6,udp6                        nfs_acl  1
  100234  1        ticotsord                                  -        superuser
  390113  1        tcp                                         -        superuser
  390435  1        tcp                                         -        superuser
  390436  1        tcp                                         -        superuser
```

```
1073741824 1 tcp,tcp6 - 1
```

The example shows information about the RPC services that are running on a server. The output that is generated by the command is filtered by the `sort` command by program number to make the information more readable. Several lines that list RPC services have been deleted from the example.

You can gather information about a particular RPC service by selecting a particular transport on a server. The following example checks the `mountd` service that is running over TCP.

```
$ rpcinfo -t bee mountd
program 100005 Version 1 ready and waiting
program 100005 Version 2 ready and waiting
program 100005 Version 3 ready and waiting
```

The following example checks the NFS service that is running over UDP.

```
$ rpcinfo -u bee nfs
program 100003 Version 2 ready and waiting
program 100003 Version 3 ready and waiting
```

snoop Command

The `snoop` command is used to monitor packets on the network. The `snoop` command must be run as the root user. The use of this command is a good way to ensure that the network hardware is functioning on both the NFS client and the NFS server.

Use the following command syntax to monitor packets on the network:

```
snoop [-d device] [-o filename] [host hostname]
```

| | |
|---------------------------------|---|
| <code>-d <i>device</i></code> | Specifies the local network interface |
| <code>-o <i>filename</i></code> | Stores all the captured packets into the named file |
| <code><i>hostname</i></code> | Displays packets going to and from a specific host only |

The `-d device` option is useful on servers that have multiple network interfaces. You can use many expressions other than setting the host. A combination of command expressions with `grep` can often generate data that is specific enough to be useful. For more information about the available options, see the [snoop\(1M\)](#) man page.

When troubleshooting, make sure that packets are going to and from the proper host. Also, look for error messages. Saving the packets to a file can simplify the review of the data.

truss Command

You can use the `truss` command to check whether a process is hung. The `truss` command must be run by the owner of the process or by `root`.

Use the following command syntax to check whether a process is hung:

```
truss [-t syscall] -p pid
```

`-t syscall` Selects system calls to trace

`-p pid` Indicates the PID of the process to be traced

`syscall` is a comma-separated list of system calls to be traced. Starting the list with an `!` character excludes the listed system calls from the trace. For more information about the available options, see the [truss\(1\)](#) man page.

EXAMPLE 13 Displaying Process Status

```
$ /usr/bin/truss -p 243
poll(0x00024D50, 2, -1)               (sleeping...)
```

The example shows that the process is waiting for another connection request, which is a normal response. If the response does not change after a new connection request has been made, the process could be hung.

For information about restarting the NFS service, see [“How to Restart NFS Service” on page 145](#). For information about troubleshooting a hung process, see [“NFS Troubleshooting Procedures” on page 140](#).

NFS Troubleshooting Procedures

To determine where the NFS service has failed, you need to follow several procedures to isolate the failure. Check for the following:

- Can the client reach the server?
- Can the client contact the NFS service on the server?
- Are the NFS services running on the server?

During this process, you might notice that other portions of the network are not functioning. For example, the name service or the physical network hardware might not be functioning.

For information about naming services, see [Working With Oracle Solaris 11.3 Directory and Naming Services: DNS and NIS](#). Also, during the process you might see that the problem is not at the client end, for example, if you receive a problem from every subnet in your work area. In this situation, assume that the problem is the server or the network hardware near the server and start the debugging process at the server, not at the client.

▼ How to Check Connectivity on an NFS Client

1. **On the client, check that the NFS server is reachable.**

```
$ /usr/sbin/ping bee  
bee is alive
```

If the command reports that the server is alive, remotely check the NFS server. For information about remotely checking the NFS server, see [“How to Check the NFS Server Remotely”](#) on page 142.

2. **If the server is not reachable from the client, ensure that the local name service is running on the client.**

For example:

- If you are using the NIS name service, verify that the `ybind` daemon is running. For more information, see [“ypbind Not Running on NIS Client”](#) in [Working With Oracle Solaris 11.3 Directory and Naming Services: DNS and NIS](#).
- If you are using the LDAP name service, verify that the `ldap_cachemgr` daemon is running. For more information, see [“Monitoring LDAP Client Status”](#) in [Working With Oracle Solaris 11.3 Directory and Naming Services: LDAP](#).

3. **If the name service is running, ensure that the client has received the correct host information.**

```
$ /usr/bin/getent hosts system
```

For example:

```
$ /usr/bin/getent hosts bee  
192.0.2.0 bee.eng.example.com
```

4. **If the host information is correct but the server is not reachable from the client, run the `ping` command from another client.**

If the command run from a second client fails, check whether the NFS service is enabled on the server. For more information, see [“How to Verify the NFS Service on the Server”](#) on page 143.

5. **If the server is reachable from the second client, use ping to check connectivity of the first client to other systems on the local network.**

If the ping command fails, check the networking software configuration on the client, for example, the `/etc/netmasks` file and the property information associated with the `svc:/system/name-service/switch` service.

6. **(Optional) Check the output of the rpcinfo command.**

If the `rpcinfo` command does not display `program 100003 Version 4 ready and waiting`, then NFS Version 4 is not enabled on the server. For information about enabling NFS Version 4, see [Table 4, “Setting Up the NFS Service,” on page 81](#).

7. **If the software is correct, check the networking hardware.**

Try to move the client to a different physical network connection.

▼ How to Check the NFS Server Remotely

Note that support for both the UDP and the MOUNT protocols is not necessary if you are using an NFS Version 4 server.

1. **Check that the NFS daemons have started on the NFS server.**

```
$ rpcinfo -s server-name | egrep 'nfs|mountd'
```

For example:

```
$ rpcinfo -s bee | egrep 'nfs|mountd'
100003 3,2 tcp,udp,tcp6,udp6 nfs superuser
100005 3,2,1 ticots,ticotsord,tcp,tcp6,ticlts,udp,udp6 mountd superuser
```

If the daemons have not been started, restart the NFS service. For more information, see [“How to Restart NFS Service” on page 145](#).

2. **On the client, test the UDP NFS connections from the server.**

```
$ /usr/bin/rpcinfo -u bee nfs
program 100003 Version 2 ready and waiting
program 100003 Version 3 ready and waiting
```

Note - NFS Version 4 does not support UDP.

If the server is running, the `rpcinfo` command lists program and version numbers that are associated with the UDP protocol. You can use the `-t` option with the `rpcinfo` command to

check the TCP connection. If the `rpcinfo` command fails, check whether the NFS service is enabled on the server. For more information, see [“How to Verify the NFS Service on the Server” on page 143](#).

3. Check that the server's `mountd` daemon is responding.

```
$ /usr/bin/rpcinfo -u bee mountd
program 100005 Version 1 ready and waiting
program 100005 Version 2 ready and waiting
program 100005 Version 3 ready and waiting
```

If the server is running, the `rpcinfo` command lists program and version numbers that are associated with the UDP protocol. Using the `-t` option tests the TCP connection. Check whether the `nfsd` and `mountd` daemon are running.

4. Check whether the local `autofs` service is used by the client by changing to a `/net` or `/home` mount point that works properly.

```
$ cd /net/eng
```

If this command fails, then as the root user on the client, restart the `autofs` service.

```
$ svcadm restart system/filesystem/autofs
```

5. Verify that the file system is shared as expected on the server.

```
$ /usr/sbin/showmount -e bee
/usr/src          eng
/export/share/man (everyone)
```

Check the entry on the server and the local mount entry for errors. Also, check the namespace. In this example, if the first client is not in the `eng` netgroup, that client cannot mount the `/usr/src` file system.

Check all entries that include mounting information in all the local files. The list includes the `/etc/vfstab` file and all the `/etc/auto_*` files.

▼ How to Verify the NFS Service on the Server

1. Become an administrator.

For more information, see [“Using Your Assigned Administrative Rights” in *Securing Users and Processes in Oracle Solaris 11.3*](#).

2. Check that the server can reach the client.

```
$ ping lilac
lilac is alive
```

3. If the client is not reachable from the server, ensure that the local name service is running on the client.
4. If the name service is running, check the networking software configuration on the server. For example, check `/etc/netmasks` and the property information associated with the `svc:/system/name-service/switch` service.
5. Check whether the `rpcbind` daemon is running on the server.

```
$ /usr/bin/rpcinfo -u localhost rpcbind
program 100000 Version 1 ready and waiting
program 100000 Version 2 ready and waiting
program 100000 Version 3 ready and waiting
```

If the server is running, the `rpcinfo` command lists program and version numbers that are associated with the UDP protocol.

6. Check whether the `nfsd` daemon is running on the server.

```
$ rpcinfo -u localhost nfs
program 100003 Version 2 ready and waiting
program 100003 Version 3 ready and waiting
$ ps -ef | grep nfsd
root 101328      0  0   Jul 12 ?          303:25 nfsd_kproc
root 101327      1  0   Jul 12 ?          2:54 /usr/lib/nfs/nfsd
root 263149 131084  0 13:59:19 pts/17    0:00 grep nfsd
```

Note - NFS Version 4 does not support UDP.

If the server is running, the `rpcinfo` command lists program and version numbers that are associated with the UDP protocol. Also, use the `-t` option with `rpcinfo` to check the TCP connection. If these commands fail, restart the NFS service. For more information, see [“How to Restart NFS Service” on page 145](#).

7. Check whether the `mountd` daemon is running on the server.

```
$ /usr/bin/rpcinfo -t localhost mountd
program 100005 Version 1 ready and waiting
program 100005 Version 2 ready and waiting
program 100005 Version 3 ready and waiting
$ ps -ef | grep mountd
root    145      1 0 Apr 07 ?          21:57 /usr/lib/autofs/automountd
```



```
root  234      1 0 Apr 07  ?    0:04  /usr/lib/nfs/mountd
root  3084 2462 1 09:30:20 pts/3  0:00  grep mountd
```

If the server is running, the `rpcinfo` command lists program and version numbers that are associated with the UDP protocol. Also, use the `-t` option with `rpcinfo` to check the TCP connection. If these commands fail, restart the NFS service. For more information, see [“How to Restart NFS Service” on page 145](#).

▼ How to Restart NFS Service

1. Become an administrator.

For more information, see [“Using Your Assigned Administrative Rights” in *Securing Users and Processes in Oracle Solaris 11.3*](#).

2. Restart the NFS service on the server.

```
$ svcadm restart network/nfs/server
```

Identifying the Host Providing the NFS Service

Use the `nfsstat` command with the `-m` option to display the current NFS information. The name of the current server is printed after `currserver=`.

For example:

```
$ nfsstat -m
/usr/local from bee,wasp:/export/share/local
Flags: vers=3,proto=tcp,sec=sys,hard,intr,llock,link,symlink,
      acl,rsize=32768,wsiz=32678,retrans=5
Failover: noresponse=0, failover=0, remap=0, currserver=bee
```

▼ How to Verify mount Command Options

No warning is issued for invalid options that are supplied with the `mount` command. This procedure helps determine whether the options that were supplied either on the command line or through the `/etc/vfstab` file were valid.

For the example in this procedure, assume that the following command has been run:

```
$ mount -F nfs -o ro,vers=2 bee:/export/share/local /mnt
```

1. Verify the options.

```
$ nfsstat -m
/mnt from bee:/export/share/local
Flags: vers=2,proto=tcp,sec=sys,hard,intr,dynamic,acl,rsize=8192,wsiz=8192,
      retrans=5
```

The file system from the bee server has been mounted with the protocol version set to 2. The `nfsstat` command does not display information about all of the options. However, using the `nfsstat` command is the most accurate way to verify the options.

2. Check the entry in the `/etc/mnttab` file.

The `mount` command does not allow invalid options to be added to the mount table. Therefore, verify that the options that are listed in the file match the options that are listed on the command line. In this way, you can check the options that are not reported by the `nfsstat` command.

```
$ grep bee /etc/mnttab
bee:/export/share/local /mnt nfs ro,vers=2,dev=2b0005e 859934818
```

Troubleshooting Autofs

Occasionally, you might encounter problems with `autofs`. This section presents a list of the error messages that `autofs` generates. The list is divided into two parts:

- Error messages that are generated by the verbose (`-v`) option of `automount`
- Error messages that might appear at any time

Each error message is followed by a description and probable cause of the message.

When troubleshooting, start the `autofs` programs with the verbose (`-v`) option.

Error Messages Generated by `automount -v`

```
bad key key in direct map mapname
```

Description: While scanning a direct map, `autofs` has found an entry key without a prefixed `/`.

Solution: Keys in direct maps must be full path names.

bad key *key* in indirect map *mapname*

Description: While scanning an indirect map, autofs has found an entry key that contains a `/`.

Solution: Indirect map keys must be simple names, not path names.

can't mount *server:pathname: reason*

Description: The mount daemon on the server refuses to provide a file handle for *server:pathname*.

Solution: Check the export table on the server.

couldn't create mount point *mountpoint: reason*

Description: Autofs was unable to create a mount point that was required for a mount. This problem most frequently occurs when you attempt to hierarchically mount all of a server's exported file systems.

Solution: A required mount point can exist only in a file system that cannot be mounted, which means the file system cannot be exported. The mount point cannot be created because the exported parent file system is exported read-only.

leading space in map entry *entry* text in *mapname*

Description: Autofs has discovered an entry in an automount map that contains leading spaces. This problem is usually an indication of an improperly continued map entry. For example:

```
fake
/blast    frobz:/usr/frotz
```

Solution: In this example, the warning is generated when autofs encounters the second line because the first line should be terminated with a backslash (`\`).

mapname: Not found

Description: The required map cannot be located. This message is produced only when the `-v` option is used.

Solution: Check the spelling and path name of the map name.

remount *server:pathname* on *mountpoint*: server not responding

Description: Autofs has failed to remount a file system that it previously unmounted.

Solution: Contact My Oracle Support for assistance. This error message is extremely rare and has no straightforward solution.

WARNING: *mountpoint* already mounted on

Description: Autofs is attempting to mount over an existing mount point. This message means that an internal error occurred in autofs (an anomaly).

Solution: Contact My Oracle Support for assistance. This error message is extremely rare and has no straightforward solution.

Miscellaneous Error Messages

dir *mountpoint* must start with '/'

Solution: The automounter mount point must be given as a full path name. Check the spelling and path name of the mount point.

hierarchical mountpoint: *pathname1* and *pathname2*

Solution: Autofs does not allow its mount points to have a hierarchical relationship. An autofs mount point must not be contained within another automounted file system.

host *server* not responding

Description: Autofs attempted to contact *server*, but received no response.

Solution: Check the NFS server status.

hostname: exports: *rpc-err*

Description: An error occurred while getting the export list from *hostname*. This message indicates a server or network problem.

Solution: Check the NFS server status.

map *mapname*, key *key*: bad

Description: The map entry is malformed and autofs cannot interpret the entry.

Solution: Recheck the entry. Perhaps the entry has characters that need to be escaped.

mapname: nis-err

Description: An error occurred when looking up an entry in a NIS map. This message can indicate NIS problems.

Solution: Check the NIS server status.

mount of *server:pathname* on *mountpoint*:reason

Description: Autofs failed to do a mount. This occurrence can indicate a server or network problem. The *reason* string defines the problem.

Solution: Contact My Oracle Support for assistance. This error message is extremely rare and has no straightforward solution.

mountpoint: Not a directory

Description: Autofs cannot mount itself on *mountpoint* because it is not a directory.

Solution: Check the spelling and path name of the mount point.

nfscast: cannot send packet: *reason*

Description: Autofs cannot send a query packet to a server in a list of replicated file system locations. The *reason* string defines the problem.

Solution: Contact My Oracle Support for assistance. This error message is extremely rare and has no straightforward solution.

nfscast: cannot receive reply: *reason*

Description: Autofs cannot receive replies from any of the servers in a list of replicated file system locations. The *reason* string defines the problem.

Solution: Contact My Oracle Support for assistance. This error message is extremely rare and has no straightforward solution.

nfscast: select: *reason*

Description: This error message indicates problem in attempting to check servers for a replicated file system. This message can indicate a network problem. The *reason* string defines the problem.

Solution: Contact My Oracle Support for assistance. This error message is extremely rare and has no straightforward solution.

pathconf: no info for *server:pathname*

Description: Autofs failed to get pathconf information for the path name.

Solution: For information about configurable path names, see the [fpathconf\(2\)](#) man page.

pathconf: *server*: server not responding

Description: Autofs is unable to contact the mount daemon on *server* that provides the information to `pathconf()`.

Solution: Avoid using the POSIX mount option with this server.

Other Errors With Autofs

If the `/etc/auto*` files have the execute bit set, the automounter tries to execute the maps, which creates messages such as the following:

```
/etc/auto_home: +auto_home: not found
```

In this situation, the `auto_home` file has incorrect permissions. Each entry in the file generates an error message that is similar to this message. Type the following command to reset the permissions to the file:

```
$ chmod 644 /etc/auto_home
```

NFS Error Messages

This section lists NFS error messages followed by a description of the conditions that can create the error and possible solutions.

Bad argument specified with index option - must be a file

Solution: You must include a file name with the `index` option. You cannot use directory names.

Cannot establish NFS service over `/dev/tcp`: transport setup problem

Description: This message is often generated when the services information in the namespace has not been updated. The message can also be reported for UDP.

Solution: To fix this problem, you must update the services data in the namespace.

For NIS and `/etc/services`, the entries must be as follows:

```
nfsd    2049/tcp    nfs    # NFS server daemon
nfsd    2049/udp    nfs    # NFS server daemon
```

Could not start *daemon*: *error*

Description: This message is displayed if the daemon terminates abnormally or if a system call error occurs. The *error* string defines the problem.

Solution: Contact My Oracle Support for assistance. This error message is rare and has no straightforward solution.

Could not use public filehandle in request to *server*

Description: This message is displayed if the `public` option is specified but the NFS server does not support the public file handle. In this situation, the mount fails.

Solution: Either try the mount request without using the public file handle or reconfigure the NFS server to support the `public` file handle.

daemon running already with pid *pid*

Description: The daemon is already running.

Solution: If you want to run a new process, terminate the current version and start a new version.

error locking *lock-file*

Description: This message is displayed when the *lock-file* that is associated with a daemon cannot be locked properly.

Solution: Contact My Oracle Support for assistance. This error message is rare and has no straightforward solution.

error checking lock-file: error

Description: This message is displayed when the *lock-file* that is associated with a daemon cannot be opened properly.

Solution: Contact My Oracle Support for assistance. This error message is rare and has no straightforward solution.

NOTICE: NFS3: failing over from *host1* to *host2*

Description: This message is displayed on the console when a failover occurs. The message is advisory only.

Solution: No action is required.

filename: File too large

Description: An NFS Version 2 client is trying to access a file that is over 2 GB.

Solution: Avoid using NFS Version 2. Mount the file system with NFS Version 3 or NFS Version 4. Also, see the description of the `nolargefiles` option in the `mount(1M)` man page.

mount: ... server not responding:RPC_PMAP_FAILURE - RPC_TIMED_OUT

Description: The server that is sharing the file system you are trying to mount is down, unreachable, or at the wrong run level, or the `rpcbind` process is dead or hung.

Solution: Wait for the server to reboot. If the server is hung, reboot the server.

mount: ... server not responding: RPC_PROG_NOT_REGISTERED

Description: The mount request is registered with the `rpcbind` process but the NFS mount daemon `mountd` is not registered.

Solution: Wait for the server to reboot. If the server is hung, reboot the server.

mount: ... No such file or directory

Description: Either the remote directory or the local directory does not exist.

Solution: Check the spelling of the directory names. Run `ls` on both directories.

mount: ...: Permission denied

Description: Your system name might not be in the list of clients or the netgroup that is allowed access to the file system you tried to mount.

Solution: Use the `showmount -e` command to verify the access list.

NFS file temporarily unavailable on the server, retrying ...

Description: An NFS Version 4 server can delegate the management of a file to a client. This message indicates that the server is recalling a delegation for another client that conflicts with a request from your client.

Solution: The recall must occur before the server can process your client's request. For more information about delegation, see [“Delegation in NFS Version 4” on page 35](#).

NFS fsstat failed for server *hostname*: RPC: Authentication error

Description: This error can be caused by many situations. One of the most difficult situations to debug is when this problem occurs because a user is in too many groups. Currently, a user can be in no more than 16 groups if the user is accessing files through NFS mounts.

Solution: An alternative does exist for users who need to be in more than 16 groups. You can use access control lists (ACLs) to provide the needed access privileges.

nfs mount: NFS can't support "nolargefiles".

Description: An NFS client has attempted to mount a file system from an NFS server by using the `-nolargefiles` option.

Solution: This option is not supported for NFS file system types.

nfs mount: NFS V2 can't support "largefiles".

Description: The NFS Version 2 protocol cannot handle large files.

Solution: You must use NFS Version 3 or NFS Version 4 if access to large files is required.

NFS server *hostname* not responding still trying

Description: If programs hang while doing file-related work, your NFS server might have failed. This message indicates that NFS server *hostname* is down or that a problem has occurred with the server or the network.

Solution: If NFS failover is being used, *hostname* is a list of servers. Start troubleshooting with [“How to Check Connectivity on an NFS Client” on page 141](#).

NFS server recovering

Description: During part of the NFS Version 4 server reboot, some operations were not permitted. This message indicates that the client is waiting for the server to permit this operation to proceed.

Solution: No action is required. Wait for the server to permit the operation.

Permission denied

Description: This message is displayed by the `ls -l`, `getfacl`, and `setfacl` commands for the following reasons:

- If the user or group that exists in an access control list (ACL) entry on an NFS Version 4 server cannot be mapped to a valid user or group on an NFS Version 4 client, the user is not allowed to read the ACL on the client.
- If the user or group that exists in an ACL entry that is being set on an NFS Version 4 client cannot be mapped to a valid user or group on an NFS Version 4 server, the user is not allowed to write or modify an ACL on the client.
- If an NFS Version 4 client and server have mismatched `nfsmapid_domain` values, ID mapping fails.

For more information about ACL entries for NFS, see [“ACLs and nfsmapid in NFS Version 4” on page 37](#).

Solution: Do the following:

- Make sure that all user IDs and group IDs in the ACL entries exist on both the client and server.
- Make sure that the value for the `nfsmapid_domain` property is set correctly in the SMF repository.

For information about the script used to determine whether any user or group cannot be mapped on the server or client, see [“Checking for Unmapped User IDs or Group IDs” on page 38](#).

port *number* in nfs URL not the same as port *number* in port option

Description: The port number that is included in the NFS URL must match the port number that is included with the `-port` option to mount. If the port numbers do not match, the mount fails.

Solution: Either change the command to make the port numbers identical or do not specify the port number that is incorrect. Usually, you do not need to specify the port number with both the NFS URL and the `-port` option.

replicas must have the same version

Description: For NFS failover to function properly, the NFS servers that are replicas must support the same version of the NFS protocol.

Solution: Do not run multiple versions.

replicated mounts must be read-only

Description: NFS failover does not work on file systems that are mounted read-write. Mounting the file system read-write increases the likelihood that a file could change.

Solution: NFS failover depends on the file systems being identical.

replicated mounts must not be soft

Description: Replicated mounts require that you wait for a timeout before NFS failover occurs.

Solution: The `soft` option requires that the mount fail immediately when a timeout starts, so you cannot include the `-soft` option with a replicated mount.

share_nfs: Cannot share more than one filesystem with 'public' option

Solution: Use the `share` command to make sure that only one file system is selected to be shared with the `-public` option. Only one public file handle can be established per server, so only one file system per server can be shared with this option.

WARNING: No network locking on `hostname:path`: contact admin to install server change

Description: An NFS client has unsuccessfully attempted to establish a connection with the Network Lock Manager on an NFS server. Rather than fail the mount, this warning is generated to warn you that locking does not work.

Solution: Upgrade the server with a newer version of the OS that provides complete lock manager support.

Accessing Network File Systems

This chapter includes a list of files and daemons that support the NFS service.

This chapter contains the following topics:

- “NFS Files” on page 157
- “NFS Daemons” on page 160

Note - If your system has zones enabled and you want to use this feature in a non-global zone, see [Introduction to Oracle Solaris Zones](#).

NFS Files

You need several files to support NFS activities on any system. Many of these files are in ASCII format, but some of the files are data files. The following table lists NFS files and their functions.

TABLE 7 NFS Files

| File Name | Function | Man Page |
|----------------------|--|---------------------------------------|
| /etc/default/fs | Specifies the default file system type for local file systems. You can determine the file system types that are supported on a client or server by checking the files in /kernel/fs. | fs(4) man page. |
| /etc/default/nfslogd | Specifies configuration information for the NFS server logging daemon, <code>nfslogd</code> . | nfslogd(1M) man page. |
| /etc/dfs/dfstab | Obsolete: Specifies the local resources to be shared. | dfstab(4) man page. |
| /etc/dfs/fstypes | Specifies the default file system types for remote file systems. The first entry defines the NFS file system type as the default. | fstypes(4) man page. |
| /etc/dfs/sharetab | Specifies the local and remote resources that are shared. <i>Do not</i> edit this file. | sharetab(4) man page. |

| File Name | Function | Man Page |
|----------------------|---|--|
| /etc/mnttab | Specifies file systems that are currently mounted, including automounted directories. Do <i>not</i> edit this file. | mnttab(4) man page. |
| /etc/netconfig | Specifies the transport protocols. Do <i>not</i> edit this file. | netconfig(4) man page. |
| /etc/nfs/nfslog.conf | Specifies general configuration information for NFS server logging. | nfslog.conf(4) man page. |
| /etc/nfs/nfslogtab | Specifies information for log post processing by the <code>nfslogd</code> daemon. Do <i>not</i> edit this file. | |
| /etc/nfssec.conf | Specifies NFS security services. | nfssec.conf(4) man page. |
| /etc/rmtab | Specifies file systems that are remotely mounted by NFS clients. Do <i>not</i> edit this file. | rmtab(4) man page. |
| /etc/vfstab | Defines file systems to be mounted locally. | vfstab(4) man page. |

/etc/default/nfslogd File

This file defines some of the parameters that are used when using NFS server logging.

Note - NFS Version 4 does not support NFS server logging.

The following parameters can be defined:

CYCLE_FREQUENCY

Determines the number of hours that must pass before the log files are cycled. The default value is 24 hours. Use this option to prevent the log files from growing too large.

IDLE_TIME

Sets the number of seconds `nfslogd` should sleep before checking for more information in the buffer file. This parameter also determines how often the configuration file is checked. This parameter, along with `MIN_PROCESSING_SIZE`, determines how often the buffer file is processed. The default value is 300 seconds. Increasing this number can improve performance by reducing the number of checks.

MAPPING_UPDATE_INTERVAL

Specifies the number of seconds between updates of the records in the file-handle-to-path mapping tables. The default value is 86400 seconds or one day. This parameter helps keep the file-handle-to-path mapping tables up to date without having to continually update the tables.

MAX_LOGS_PRESERVE

Determines the number of log files to be saved. The default value is 10.

MIN_PROCESSING_SIZE

Sets the minimum number of bytes that the buffer file must reach before processing and writing to the log file. This parameter, along with `IDLE_TIME`, determines how often the buffer file is processed. The default value is 524288 bytes. Increasing this number can improve performance by reducing the number of times the buffer file is processed.

PRUNE_TIMEOUT

Selects the number of hours that must pass before a file-handle-to-path mapping record times out and can be reduced. The default value is 168 hours or 7 days.

UMASK

Specifies the file mode creation mask for the log files that are created by `nfslogd`. The default value is 0137.

`/etc/nfs/nfslog.conf` File

This file defines the path, file names, and type of logging to be used by `nfslogd`. Each definition is associated with a tag. Starting NFS server logging requires that you identify the tag for each file system. The global tag defines the default values.

Note - NFS Version 4 does not support NFS server logging.

You can use the following parameters with each tag as needed.

`defaultdir=path`

Specifies the default directory path for the logging files. Unless you specify differently, the default directory is `/var/nfs`.

`log=path/filename`

Sets the path and file name for the log files. The default is `/var/nfs/nfslog`.

`fhtable=path/filename`

Selects the path and file name for the file-handle-to-path database files. The default is `/var/nfs/fhtable`.

`buffer=path/filename`

Determines the path and file name for the buffer files. The default is `/var/nfs/nfslog_workbuffer`.

`logformat=basic|extended`

Selects the format to be used when creating user-readable log files. The basic format produces a log file that is similar to some ftpd daemons. The extended format gives a more detailed view.

If the path is not specified, the path that is defined by `defaultdir` is used. Also, you can override `defaultdir` by using an absolute path.

To identify the files more easily, place the files in separate directories. The following example shows the changes that are needed.

EXAMPLE 14 Sample NFS Server Logging Configuration File

```
$ cat /etc/nfs/nfslog.conf
#ident "@(#)nfslog.conf 1.5 99/02/21 SMI"
#
.
.
# NFS server log configuration file.
#

global defaultdir=/var/nfs \
        log=nfslog fhtable=fhtable buffer=nfslog_workbuffer

publicftp log=logs/nfslog fhtable=fh/fhtables buffer=buffers/workbuffer
```

The example shows a file system that is shared with `log=publicftp`. A file system that is shared with `log=publicftp` uses the following values:

- The default directory is `/var/nfs`.
- Log files are stored in `/var/nfs/logs/nfslog*`.
- File-handle-to-path database tables are stored in `/var/nfs/fh/fhtables`.
- Buffer files are stored in `/var/nfs/buffers/workbuffer`.

For information about enabling NFS server logging, refer to [“How to Enable NFS Server Logging” on page 73](#).

NFS Daemons

To support NFS activities, several daemons are started when a system goes into run level or multiuser mode. The `mountd` and `nfsd` daemons are run on systems that are servers. The

automatic startup of the server daemons depends on the existence of at least one NFS share. To display the current list of NFS shares, run the `share -F nfs` command. To support NFS file locking, the `lockd` and `statd` daemons are run on NFS clients and servers. However, unlike previous versions of NFS, in NFS Version 4, the daemons `lockd`, `statd`, and `nfslogd` are not used.

This section describes the following daemons.

- [“automountd Daemon” on page 161](#)
- [“lockd Daemon” on page 162](#)
- [“mountd Daemon” on page 163](#)
- [“nfs4cbd Daemon” on page 164](#)
- [“nfsd Daemon” on page 164](#)
- [“nfslogd Daemon” on page 165](#)
- [“nfsmapid Daemon” on page 165](#)
- [“reparseid Daemon” on page 172](#)
- [“statd Daemon” on page 172](#)

automountd Daemon

The `automountd` daemon handles the mounting and unmounting requests from the `autofs` service. The syntax of the command is as follows:

```
automountd [-Tnv] [-D name=value]
```

| | |
|----------------------------|--|
| <code>-T</code> | Enables tracing. |
| <code>-n</code> | Disables browsing on all <code>autofs</code> nodes. |
| <code>-v</code> | Logs all status messages to the console. |
| <code>-D name=value</code> | Substitutes <i>value</i> for the automount map variable that is indicated by <i>name</i> . |

The default value for the automount map is `/etc/auto_master`. Use the `-T` option for troubleshooting.

You can make the same specifications with the `sharectl` command that you would make on the command line. However, unlike the command-line options, the SMF repository preserves your specifications, through service restarts and system reboots, as well as system upgrades. You can set the following parameters for the `automountd` daemon.

automountd_verbose

Logs status messages to the console and is the equivalent of the `-v` argument for the `automountd` daemon. The default value is `FALSE`.

nobrowse

Turns browsing on or off for all `autofs` mount points and is the equivalent of the `-n` argument for `automountd`. The default value is `FALSE`.

trace

Expands each remote procedure call (RPC) and displays the expanded RPC on standard output. This keyword is the equivalent of the `-T` argument for `automountd`. The default value is 0. Values can range from 0 to 5.

environment

Permits you to assign different values to different environments. This keyword is the equivalent of the `-D` argument for `automountd`. The `environment` parameter can be used multiple times. However, you must use separate entries for each environment assignment.

lockd Daemon

The `lockd` daemon supports record-locking operations on NFS files. The `lockd` daemon manages RPC connections between the client and the server for the Network Lock Manager (NLM) protocol. The daemon is normally started without any options. You can use three options with this command. You can set these options either from the command line or by setting parameters using the `sharectl` command. For more information, see the [lockd\(1M\)](#) man page.

Note - The `LOCKD_GRACE_PERIOD` keyword and the `-g` option have been deprecated. The deprecated keyword is replaced with the new `grace_period` parameter. If both keywords are set, the value for `grace_period` overrides the value for `LOCKD_GRACE_PERIOD`.

Like `LOCKD_GRACE_PERIOD`, the `grace_period=graceperiod` parameter sets the number of seconds after a server reboot that the clients have to reclaim both NFS Version 3 locks, provided by NLM, and NFS Version 4 locks.

The `lockd_retransmit_timeout=timeout` parameter selects the number of seconds to wait before retransmitting a lock request to the remote server. This option affects the NFS client-side service. The default value for `timeout` is 5 seconds. Decreasing the `timeout` value can improve response time for NFS clients on a "noisy" network. However, this change can cause additional

server load by increasing the frequency of lock requests. The same parameter can be used from the command line by starting the daemon with the `-t timeout` option.

The `lockd_servers=number` parameter specifies the maximum number of concurrent lockd requests. The default value is 1024.

The `nthreads` parameter specifies the maximum number of concurrent threads that the server can handle. All NFS clients that use UDP share a single connection with the NFS server. Under these conditions, you might have to increase the number of threads that are available for the UDP connection. A minimum calculation would be to allow two threads for each UDP client. However, this number is specific to the workload on the client, so two threads per client might not be sufficient. The disadvantage to using more threads is that when the threads are used, more memory is used on the NFS server. If the threads are never used, however, increasing `nthreads` has no effect. The same parameter can be used from the command line by starting the daemon with the `nthreads` option.

mountd Daemon

The `mountd` daemon handles file system mount requests from remote systems and provides access control. The `mountd` daemon checks `/etc/dfs/sharetab` to determine which file systems are available for remote mounting and which systems are allowed to do the remote mounting. For more information, see the [mountd\(1M\)](#) man page.

- `-v` Runs the command in verbose mode. Every time an NFS server determines the access that a client should be granted, a message is printed on the console. The information that is generated can be useful when trying to determine why a client cannot access a file system.
- `-r` Rejects all future mount requests from clients. This option does not affect clients that already have a file system mounted.

In addition to the command-line options, several SMF parameters can be used to configure the `mountd` daemon:

`client_versmin`

Sets the minimum version of the NFS protocol to be used by the NFS client. The default is 2. Other valid values include 3 or 4. Refer to [“Setting Up the NFS Service” on page 81](#).

`client_versmax`

Sets the maximum version of the NFS protocol to be used by the NFS client. The default is 4. Other valid values include 2 or 3. Refer to [“Setting Up the NFS Service” on page 81](#).

nfs4cbd Daemon

The `nfs4cbd` daemon, which is for the exclusive use of the NFS Version 4 client, manages the communication endpoints for the NFS Version 4 callback program. The daemon has no user-accessible interface. For more information, see the [nfs4cbd\(1M\)](#) man page.

nfsd Daemon

The `nfsd` daemon handles client file system requests. You can use several options with this command. See the [nfsd\(1M\)](#) man page for a complete listing. These options can either be used from the command line or by setting the appropriate SMF parameter with the `sharectl` command.

`listen_backlog=length` Sets the length of the connection queue over connection-oriented transports for NFS and TCP. The default value is 32 entries. The same selection can be made from the command line by starting `nfsd` with the `-l` option.

`max_connections=#-conn` Selects the maximum number of connections per connection-oriented transport. The default value for `#-conn` is unlimited. The same parameter can be used from the command line by starting the daemon with the `-c #-conn` option.

`servers=nservers` Selects the maximum number of concurrent requests that a server can handle. The default value for `nservers` is 1024. The same selection can be made from the command line by starting `nfsd` with the `nservers` option.

Unlike older versions of this daemon, `nfsd` does not spawn multiple copies to handle concurrent requests. Checking the process table with `ps` only shows one copy of the daemon running.

In addition, the following SMF parameters can be used to configure the `mountd` daemon. These parameters do not have command-line equivalents:

`server_versmin`

Sets the minimum version of the NFS protocol to be registered and offered by the server. The default is 2. Other valid values include 3 or 4. Refer to [“Setting Up the NFS Service” on page 81](#).

`server_versmax`

Sets the maximum version of the NFS protocol to be registered and offered by the server. The default is 4. Other valid values include 2 or 3. Refer to [“Setting Up the NFS Service” on page 81](#).

`server_delegation`

Controls whether the NFS Version 4 delegation feature is enabled for the server. If this feature is enabled, the server attempts to provide delegations to the NFS Version 4 client. By default, server delegation is enabled. To disable server delegation, see [“How to Select Different Versions of NFS on a Server” on page 82](#). For more information, refer to [“Delegation in NFS Version 4” on page 35](#).

nfslogd Daemon

Note - NFS Version 4 does not use this daemon.

The `nfslogd` daemon provides operational logging. NFS operations that are logged against a server are based on the configuration options that are defined in `/etc/default/nfslogd`. When NFS server logging is enabled, records of all RPC operations on a selected file system are written to a buffer file by the kernel. Then `nfslogd` postprocesses these requests. The name service switch is used to help map UIDs to logins and IP addresses to host names. The number is recorded if no match can be found through the identified name services.

Mapping of file handles to path names is also handled by `nfslogd`. The daemon tracks these mappings in a file-handle-to-path mapping table. One mapping table exists for each tag that is identified in `/etc/nfs/nfslogd`. After post-processing, the records are written to ASCII log files.

nfsmapid Daemon

Version 4 of the NFS protocol (RFC3530) changed the way user or group identifiers (UID or GID) are exchanged between the client and server. The protocol requires that a file's owner and group attributes be exchanged between an NFS Version 4 client and an NFS Version 4 server as strings in the form `user@nfsv4-domain` or `group@nfsv4-domain`, respectively.

For example, user `known_user` has a UID 123456 on an NFS Version 4 client whose fully qualified hostname is `system.example.com`. For the client to make requests to the NFS Version

4 server, the client must map the UID 123456 to `known_user@example.com` and then send this attribute to the NFS Version 4 server. After the server receives `known_user@example.com` from the client, the server maps the string to the local UID 123456, which is understood by the underlying file system. This functionality assumes that every UID and GID in the network is unique and that the NFS Version 4 domains on the client match the NFS Version 4 domains on the server.

The NFS Version 4 client and server are both capable of performing integer-to-string and string-to-integer conversions. For example, in response to a `GETATTR` operation, the NFS Version 4 server maps UIDs and GIDs obtained from the underlying file system into their respective string representation and sends this information to the client. Alternately, the client must also map UIDs and GIDs into string representations. For example, in response to the `chown` command, the client maps the new UID or GID to a string representation before sending a `SETATTR` operation to the server.

Note, however, that the client and server respond differently to unrecognized strings:

- If the user does not exist on the server, even within the same NFS Version 4 domain configuration, the server rejects the remote procedure call (RPC) and returns an error message to the client. This situation limits the operations that can be performed by the remote user.
- If the user exists on both the client and server but they have mismatched domains, the server rejects the attribute modifying operations (such as `SETATTR`) that require the server to map the inbound user string to an integer value that the underlying file system can understand. For NFS Version 4 clients and servers to function properly, their NFS Version 4 domains, the portion of the string after the `@` sign, should match.
- If the NFS Version 4 client does not recognize a user or group name from the server, the client is unable to map the string to its unique ID, an integer value. Under such circumstances, the client maps the inbound user or group string to the `nobody` user. This mapping to `nobody` creates varied problems for different applications. With NFS Version 4, operations that modify file attributes will fail.
- If the server does not recognize the given user or group name, even if the NFS Version 4 domains match, the server is unable to map the user or group name to its unique ID, an integer value. Under such circumstances, the server maps the inbound user or group name to the `nobody` user. To prevent such occurrences, administrators should avoid making special accounts that exist only on the NFS Version 4 client.

You can change the domain name for the clients and servers using the `sharectl` command with the `nfsmapid_domain` option. This option sets a common domain for clients and servers. It overrides the default behavior of using the local DNS domain name. For task information, refer to [“Setting Up the NFS Service” on page 81](#).

Configuration Files and the `nfsmapid` Daemon

The `nfsmapid` daemon uses the SMF configuration information found in `svc:system/name-service/switch` and in `svc:/network/dns/client` as follows:

- `nfsmapid` uses standard C library functions to request password and group information from back-end name services. These name services are controlled by the settings in the `svc:system/name-service/switch` SMF service. Any changes to the service properties affect `nfsmapid` operations. For more information about the `svc:system/name-service/switch` SMF service, see the [`nsswitch.conf\(4\)`](#) man page.
- To ensure that NFS Version 4 clients are capable of mounting file systems from different domains, `nfsmapid` relies on the configuration of the DNS TXT resource record (RR), `_nfsv4idmapdomain`. For more information about configuring the `_nfsv4idmapdomain` resource record, see “[`nfsmapid` and DNS TXT Records](#)” on page 168. Also, note the following:
 - The DNS TXT RR should be explicitly configured on the DNS server with the desired domain information.
 - The `svc:system/name-service/switch` SMF service should be configured to enable the resolver to find the DNS server and search the TXT records for client and server NFS Version 4 domains.

For more information, see the following:

- “[Precedence Rules](#)” on page 167
- “[Configuring the NFS Version 4 Default Domain](#)” on page 170
- The [`resolv.conf\(4\)`](#) man page

Precedence Rules

For `nfsmapid` to work properly, NFS Version 4 clients and servers must have the same domain. To ensure matching NFS Version 4 domains, `nfsmapid` follows these strict precedence rules:

1. The daemon first checks the SMF repository for a value that has been assigned to the `nfsmapid_domain` parameter. If a value is found, the assigned value takes precedence over any other settings. The assigned value is appended to the outbound attribute strings and is compared against inbound attribute strings. For procedural information, see “[Setting Up the NFS Service](#)” on page 81.

Note - The use of the NFSMAPID_DOMAIN setting is not scalable and is not recommended for large deployments.

2. If no value has been assigned to `nfsmapid_domain`, then the daemon checks for a domain name from a DNS TXT RR. `nfsmapid` relies on directives in the `/etc/resolv.conf` file that are used by the set of routines in the `resolver`. The `resolver` searches through the configured DNS servers for the `_nfsv4idmapdomain` TXT RR. Note that the use of DNS TXT records is more scalable. For this reason, continued use of TXT records is much preferred over setting the parameter in the SMF repository.
3. If no DNS TXT record is configured to provide a domain name, then the `nfsmapid` daemon uses the value specified by the `domain` or `search` directive in the `/etc/resolv.conf` file, with the directive specified last taking precedence.

In the following example, both the `domain` and `search` directives are used. The `nfsmapid` daemon uses the first domain listed after the `search` directive, which is `example.com`.

```
domain company.example.com
search example.com corp.example.com
```

4. If the `/etc/resolv.conf` file does not exist, `nfsmapid` obtains the NFS Version 4 domain name by following the behavior of the `domainname` command. Specifically, if the `/etc/defaultdomain` file exists, `nfsmapid` uses the contents of that file for the NFS Version 4 domain. If the `/etc/defaultdomain` file does not exist, `nfsmapid` uses the domain name that is provided by the network's configured naming service. For more information, see the [domainname\(1M\)](#) man page.

nfsmapid and DNS TXT Records

The ubiquitous nature of DNS provides an efficient storage and distribution mechanism for the NFS Version 4 domain name. Additionally, because of the inherent scalability of DNS, the use of DNS TXT resource records is the preferred method for configuring the NFS Version 4 domain name for large deployments. You should configure the `_nfsv4idmapdomain` TXT record on enterprise-level DNS servers. Such configurations ensure that any NFS Version 4 client or server can find its NFS Version 4 domain by traversing the DNS tree.

The following example shows a preferred entry for enabling the DNS server to provide the NFS Version 4 domain name:

```
_nfsv4idmapdomain IN TXT "corp.example"
```


In this example, the domain name to configure is the value that is enclosed in double quotes. Note that no `t11` field is specified and that no domain is appended to `_nfsv4idmapdomain`, which is the value in the `owner` field. This configuration enables the TXT record to use the zone's `$_ORIGIN` entry from the Start-Of-Authority (SOA) record. For example, at different levels of the domain namespace, the record could read as follows:

```
_nfsv4idmapdomain.subnet.example.com.    IN    TXT    "corp.example"
_nfsv4idmapdomain.example.com.          IN    TXT    "corp.example"
```

This configuration provides DNS clients with the added flexibility of using the `resolv.conf` file to search up the DNS tree hierarchy. See the [resolv.conf\(4\)](#) man page. This capability provides a higher probability of finding the TXT record. For even more flexibility, lower level DNS subdomains can define their own DNS TXT resource records (RRs). This capability enables lower level DNS subdomains to override the TXT record that is defined by the top level DNS domain.

Note - The domain that is specified by the TXT record can be an arbitrary string that does not necessarily match the DNS domain for clients and servers that use NFS Version 4. You have the option of not sharing NFS Version 4 data with other DNS domains.

Checking for the NFS Version 4 Domain

Before assigning a value for your network's NFS Version 4 domain, check whether an NFS Version 4 domain has already been configured for your network. The following examples provide ways of identifying your network's NFS Version 4 domain.

- To identify the NFS Version 4 domain from a DNS TXT RR, use either the `nslookup` or the `dig` command:

The following example shows sample output for the `nslookup` command:

```
$ nslookup -q=txt _nfsv4idmapdomain
Server:          10.255.255.255
Address:         10.255.255.255#53

_nfsv4idmapdomain.company.example.com text = "example.com"
```

The following example shows sample output for the `dig` command:

```
$ dig +domain=company.example.com -t TXT _nfsv4idmapdomain
...
;; QUESTION SECTION:
```

```

;_nfsv4idmapdomain.company.example.com. IN    TXT

;; ANSWER SECTION:
_nfsv4idmapdomain.company.example.com. 21600 IN TXT    "example.com"

;; AUTHORITY SECTION:
...
```

For information about setting up a DNS TXT RR, see [“nfsmapid and DNS TXT Records” on page 168](#).

- If your network is not set up with a NFS Version 4 DNS TXT RR, use the following command to identify your NFS Version 4 domain from the DNS domain name:

```
$ egrep domain /etc/resolv.conf
domain company.example.com
```

- If the `/etc/resolv.conf` file is not configured to provide a DNS domain name for the client, use the following command to identify the domain from the network's NFS Version 4 domain configuration:

```
$ cat /system/volatile/nfs4_domain
example.com
```

- If you are using a different naming service, such as NIS, use the following command to identify the domain for the naming service configured for your network:

```
$ domainname
it.company.example.com
```

For more information, see the following man pages:

- [nslookup\(1M\)](#)
- [dig\(1M\)](#)
- [resolv.conf\(4\)](#)
- [domainname\(1M\)](#)

Configuring the NFS Version 4 Default Domain

This section describes how the network obtains the desired default domain:

- For most current releases, see [“Configuring an NFS Version 4 Default Domain in the Oracle Solaris 11 Release” on page 171](#).
- For the initial Oracle Solaris 10 release, see [“Configuring an NFS Version 4 Default Domain in the Oracle Solaris 10 Release” on page 171](#).

Configuring an NFS Version 4 Default Domain in the Oracle Solaris 11 Release

In the Oracle Solaris 11 release, set the default NFS domain version by typing the following command:

```
$ sharectl set -p nfsmapid_domain=example.com nfs
```

Note - Because of the inherent ubiquitous and scalable nature of DNS, the use of DNS TXT records for configuring the domain of large NFS Version 4 deployments continues to be preferred and strongly encouraged. See [“nfsmapid and DNS TXT Records” on page 168](#).

Configuring an NFS Version 4 Default Domain in the Oracle Solaris 10 Release

In the initial Oracle Solaris 10 release of NFS Version 4, if your network included multiple DNS domains but only had a single UID and GID namespace, all clients had to use one value for `nfsmapid_domain`. For sites that use DNS, `nfsmapid` resolves this issue by obtaining the domain name from the value that you assigned to `_nfsv4idmapdomain`. For more information, see [“nfsmapid and DNS TXT Records” on page 168](#). If your network is not configured to use DNS, during the first system boot, the operating system uses the `sysidconfig` utility to provide the following prompts for an NFS Version 4 domain name:

```
This system is configured with NFS Version 4, which uses a
domain name that is automatically derived from the system's
name services. The derived domain name is sufficient for most
configurations. In a few cases, mounts that cross different
domains might cause files to be owned by nobody due to the
lack of a common domain name.
```

```
Do you need to override the system's default NFS version 4 domain
name (yes/no)? [no]
```

The default response is [no]. If you choose [no], you see the following message:

```
For more information about how the NFS Version 4 default domain name is
derived and its impact, refer to the man pages for nfsmapid(1M) and
nfs(4), and the System Administration Guide: Network Services.
```

If you choose [yes], you see this prompt:

```
Enter the domain to be used as the NFS Version 4 domain name.
NFS Version 4 domain name []:
```

Note - If a value for `nfsmapid_domain` exists in the SMF repository, the domain name that you provide overrides that value.

Additional Information About `nfsmapid`

For more information about `nfsmapid`, see the following:

- [nfsmapid\(1M\)](#) man page
- [nfs\(4\)](#) man page
- [RFC 1464, Using the Domain Name System To Store Arbitrary String Attributes](#)
- “ACLs and `nfsmapid` in NFS Version 4” on page 37

reparse Daemon

The `reparse` daemon interprets the data associated with a reparse point. These points are used by DFS and NFS referrals on SMB and NFS file servers. This service is managed by SMF and should not be manually started.

statd Daemon

Note - NFS Version 4 does not use this daemon.

The `statd` daemon works with `lockd` to provide crash and recovery functions for the lock manager. The `statd` daemon tracks the clients that hold locks on an NFS server. If a server crashes, on rebooting, `statd` on the server contacts `statd` on the client. The client `statd` can then attempt to reclaim any locks on the server. The client `statd` also informs the server `statd` when a client has crashed so that the client's locks on the server can be cleared. This daemon has no options. For more information, see the [statd\(1M\)](#) man page.

NFS File Sharing Command Reference

Administering Network File Systems

The following list shows the basic command syntax for performing some common network file sharing tasks in the Oracle Solaris 11.3 release. For more information, see the [nfsref\(1M\)](#), [nfsstat\(1M\)](#), [share\(1M\)](#), and [mount\(1M\)](#) man pages.

- Share a ZFS file system
 - \$ **zfs set share.nfs=on | off** *filesystem*
- Define file systems to use NFS server logging
 - \$ **share -F nfs -o,log=global** *filesystem*
- List the shared NFS file systems
 - \$ **share -F nfs**
- Mount a file system
 - \$ **mount -F nfs -o specific-options** *resource mount-point*
- Mount an NFS File System by Using an NFS URL
 - \$ **mount -F nfs nfs://host[:port]/pathname** *mount-point*
- Display information about file systems available for mounting
 - \$ **/usr/sbin/showmount -e** *hostname*
- Select different versions of NFS on a server
 - \$ **sharectl set -p server_versmax=version_num** **nfs**
 - \$ **sharectl set -p server_versmin=version_num** **nfs**
- Disable server delegation
 - \$ **sharectl set -p server_delegation=off** **nfs**
- Set a common domain for NFS clients and servers

- \$ **sharectl set -p nfsmapiid_domain=domain_name nfs**
- Create an NFS referral
 - \$ **nfsref add ref_filesystem server_name1:filesystem server_name2:filesystem**
- Remove an NFS referral
 - \$ **nfsref remove ref_filesystem**
- Disable autofs browsability on a single NFS client
 - \$ **sharectl set -p nobrowse=TRUE autofs**

Troubleshooting Network File Systems

The following list shows the basic command syntax for some common NFS troubleshooting tasks in the Oracle Solaris 11.3 release. For more information, see the [nfsref\(1M\)](#) and [nfsstat\(1M\)](#) man pages.

- Remove all file, record, and share locks for an NFS client
 - \$ **clear_locks -s hostname**
- Display file system operations statistics by file system type and by mount point
 - \$ **fsstat nfsversnum nfsversnum filesystem**
- Display NFS server statistics
 - \$ **nfsstat -s**
- Display the current NFS information
 - \$ **nfsstat -m**
- Display stack trace for NFS process
 - \$ **/usr/bin/pstack nfs_process_id**
- Display RPC service information for a host
 - \$ **rpcinfo -s hostname**

Index

Numbers and Symbols

- # (pound sign)
 - comments in direct maps, 56
 - comments in indirect maps, 58
 - comments in master map (auto_master), 54
- & (ampersand)
 - in autofs maps, 69
- * (asterisk)
 - in autofs maps, 70
- + (plus sign)
 - in autofs map names, 66, 67
- (dash)
 - in autofs map names, 66
- / (slash)
 - /- as master map mount point, 54, 57
 - master map names preceded by, 54
 - root directory
 - mounting by diskless clients, 14
- \ (backslash) in autofs maps, 54, 56, 58

A

- a option
 - showmount command, 131
 - umount command, 120
- access control list (ACL) and NFS
 - description, 18
 - error message, Permission denied, 154
- access control lists (ACLs) and NFS
 - description, 37
- accessing
 - NFS referrals, 90
- already mounted message, 148

- ampersand (&)
 - in autofs maps, 69
- anon option
 - share command, 126
- applications
 - hung, 154
- ARCH map variable, 65
- asterisk (*)
 - in autofs maps, 70
- authentication
 - DH, 47
 - Diffie-Hellman (DH), 47
 - RPC, 47
 - UNIX, 45, 47
- auto_home map
 - /home directory, 100
 - /home directory server setup, 101
 - /home mount point, 53, 55
- auto_master file
 - nobrowse option, 108
- autofs
 - /home directory, 100
 - administering maps, 97
 - browsability, 22, 107
 - consolidating project-related files, 102
 - features, 22
 - home directory server setup, 101
 - maps
 - browsability and, 22
 - direct, 56, 57
 - indirect, 57, 59
 - master, 53, 54
 - network navigation, 53

- read-only file selection, 61, 64
 - referring to other maps, 66, 67
 - starting the navigation process, 55, 59
 - types of, 97
 - variables, 65, 65
 - metacharacters, 69, 70
 - mount process, 60, 61
 - mounting file systems, 77
 - namespace data, 22
 - NFS URL and, 107
 - nobrowse option, 108
 - non-NFS file system access, 100
 - overview, 14
 - public file handle and, 107
 - reference information for, 69, 70
 - replicating shared files across several servers, 106
 - shared namespace access, 104
 - special characters, 70
 - supporting incompatible client OS versions using, 105
 - troubleshooting, 146
 - unmounting process, 61
 - automount command, 112
 - autofs and, 14
 - error messages, 146
 - modifying autofs master map (auto_master), 98
 - overview, 51
 - v option, 146
 - when to run, 98
 - automountd daemon, 161
 - autofs and, 14
 - description, 22
 - mounting and, 22
 - overview, 51, 52
 - avoiding problems with ACLs in NFS, 38
- B**
- background file-mounting option, 115
 - backslash (\) in maps, 54, 56, 58
 - bad argument specified with index option, 150
 - bad key message, 147
 - bg option
 - mount command, 115
 - booting
 - diskless client security, 48
 - mounting file systems, 75
 - browsability
 - disabling, 107
 - overview, 22
 - browsing
 - with an NFS URL, 89
- C**
- cache and NFS Version 3, 16
 - can't mount message, 147
 - cannot receive reply message, 149
 - cannot send packet message, 149
 - checking for unmapped user or group IDs, 38
 - clear_locks command, 113
 - client recovery
 - NFS Version 4, 33
 - client-side failover
 - enabling, 78
 - in NFS Version 4, 42
 - NFS locking and, 42
 - NFS support, 20
 - overview, 41
 - replicated file systems, 42
 - client_versmax parameter, 163
 - client_versmin parameter, 163
 - commands
 - FedFS, 132
 - hung programs, 154
 - NFS, 111
 - comments
 - in direct maps, 56
 - in indirect maps, 58
 - in master map (auto_master), 54
 - consolidating project-related files, 102
 - conversation key, 47
 - could not use public filehandle message, 151
 - couldn't create mount point message, 147
 - CPU map variable, 65
 - creating

- namespace database (FedFS), 92
 - NFS referrals, 50, 90, 94
 - secure connection (FedFS), 93
 - credentials
 - description, 46
 - UNIX authentication, 47
- D**
- d option
 - showmount command, 131
 - daemon running already message, 151
 - daemons
 - automountd, 161
 - autofs and, 14
 - overview, 51, 52
 - lockd, 162
 - mountd, 163
 - checking response on server, 143
 - not registered with rpcbind daemon, 152
 - verifying if running, 144, 152
 - nfs4cbd, 164
 - nfsd
 - checking response on server, 142
 - description, 164
 - verifying if running, 144
 - nfslogd, 165
 - nfsmapid, 165
 - repared, 172
 - required for remote mounting, 133
 - rpcbind
 - mount error messages, 152, 152
 - statd, 172
 - dash (-)
 - in autofs map names, 66
 - delegation
 - NFS Version 4, 35
 - DH authentication
 - overview, 47, 47
 - password protection, 46
 - secure NFS and, 85
 - user authentication, 45
 - dir must start with '/' message, 148
 - direct I/O mounting option, 115
 - direct maps (autofs)
 - comments in, 56
 - description, 98
 - example, 56
 - overview, 57
 - syntax, 56
 - when to run automount command, 98
 - disabling
 - autofs browsability
 - overview, 107
 - mount access for one client, 78
 - diskless clients
 - manual mounting requirements, 14
 - security during boot process, 48
 - displaying
 - mountable file systems, 80
 - restricted file system information, 80
 - DNS record
 - FedFS, 91
 - domain names
 - Secure NFS system and, 85
 - domains
 - definition, 85
- E**
- /etc/default/autofs file
 - configuring autofs environment, 97
 - /etc/default/nfslogd file, 158
 - /etc/mnttab file
 - comparing with auto_master map, 52
 - /etc/netconfig file, 158
 - /etc/nfs/nfslog.conf file, 159
 - /etc/services file
 - nfsd entries, 151
 - /etc/vfstab file
 - automount command and, 53
 - diskless clients, 14
 - enabling client-side failover, 78
 - mounting file systems at boot time, 75

- NFS servers and, 75
- e option
 - showmount command, 131
- enabling
 - client-side failover, 78
 - NFS server logging, 73
 - WebNFS service, 88
- error checking message, 152
- error locking message, 151
- error messages
 - generated by automount -v, 146
 - miscellaneous automount messages, 148
 - No such file or directory, 152
 - open errors
 - NFS and, 16
 - Permission denied, 153
 - server not responding
 - hung programs, 154
 - keyboard interrupt for, 133
 - remote mounting problems, 152, 154
 - write errors
 - NFS and, 16
- executable maps, 67

F

- F option
 - unshareall command, 130
- failover
 - error message, 152
 - mount command example, 118, 118
 - NFS support, 20
- Federated File System *See* FedFS
- FedFS
 - administering, 91
 - DNS record for, 91
 - LDAP schema, 92
 - mount point, 55
 - mounting, 91
- FedFS commands, 132
- fg option
 - mount command, 115

- file attributes and NFS Version 3, 16
- file permissions
 - NFS Version 3 improvement, 16
 - WebNFS and, 88
- File Sharing
 - NFS Commands, 173
- file sharing
 - automatic, 72
 - examples, 128
 - giving root access, 127
 - listed clients only, 125
 - multiple file systems, 130
 - NFS Version 3 improvements, 16, 19
 - overview, 124
 - read-only access, 125, 125, 128
 - read-write access, 125, 128
 - replicating shared files across several servers, 106
 - security issues, 45, 125, 127
 - unauthenticated users and, 126
 - unsharing, 130
- file system namespace
 - NFS Version 4, 29
- file systems and NFS, 16
- file too large message, 152
- file transfer size
 - negotiation, 39
- file-sharing options, 125
- files and file systems
 - autofs access
 - non-NFS file systems, 100
 - autofs selection of files, 61, 64
 - consolidating project-related files, 102
 - definition, 16
 - file systems, 16
 - local file systems
 - unmounting groups, 121
 - NFS files and their functions, 157
 - NFS treatment of, 16, 16
 - remote file systems
 - listing clients with remotely mounted file systems, 131
 - mounting from file system table, 121
 - unmounting groups, 121

firewalls
 mounting file systems through, 79
 NFS access through, 21
 WebNFS access through, 90

forcedirectio option
 mount command, 115

foreground file-mounting option, 115

ftp archive
 WebNFS and, 88

fuser command
 umountall command and, 121

G

-g option
 lockd daemon, 162

grace_period parameter
 lockd daemon, 162

GSS-API
 and NFS, 21

H

/home directory and NFS server setup, 101

/home mount point, 53, 55

-h option
 umountall command, 121

hard option
 mount command, 117

hierarchical mount points message, 148

hierarchical mounts (multiple mounts), 60

HOST map variable, 65

host not responding message, 148

hosts
 unmounting all file systems from, 121

HTML file
 WebNFS and, 88

httpd command
 firewall access and WebNFS, 90

hung programs, 154

I

ID mapping fails
 reasons why, 37

index option
 bad argument error message, 150
 share command, 89
 WebNFS and, 88

indirect maps (autofs)
 comments in, 58
 description, 98
 example, 58, 59
 overview, 57, 59
 syntax, 57, 58
 when to run automount command, 98

-intr option
 mount command, 133

K

/kernel/fs file
 checking, 157

-k option
 umountall command, 121

KERB authentication
 NFS and, 20

kernel
 checking response on server, 141

keyboard interruption of mounting, 133

keylogin command
 remote login security issues, 48

keylogout command
 secure NFS and, 48

keywords
 NFS Version negotiation, 28

L

-l option
 umountall command, 121

large files
 NFS support, 20

largefiles option

- error message, 153
- mount command, 115
- LDAP schema
 - for FedFS, 92
- leading space in map entry message, 147
- listing
 - clients with remotely mounted file systems, 131
 - mounted file systems, 119
 - shared file systems, 128
- local cache and NFS Version 3, 16
- local file systems
 - unmounting groups, 121
- local files
 - updating autofs maps, 98
- lockd daemon, 162
- LOCKD_GRACE_PERIOD parameter
 - lockd daemon, 162
- lockd_retransmit_timeout parameter
 - lockd daemon, 162
- lockd_servers parameter
 - lockd daemon, 163
- locking
 - NFS Version 3 improvements, 19
- log option
 - share command, 126
- login command
 - secure NFS and, 48
- ls command
 - ACL entries and, 37

- M**
- map key bad message, 149
- maps (autofs)
 - administrative tasks, 97
 - automount command
 - when to run, 98
 - avoiding mount conflicts, 99
 - comments in, 54, 56, 58
 - direct, 56, 57
 - executable, 67
 - indirect, 57, 59
 - maintenance methods, 98
 - master, 53, 54
 - multiple mounts, 60
 - network navigation, 53
 - referring to other maps, 66, 67
 - selecting read-only files for clients, 61, 64
 - special characters in, 70
 - splitting long lines in, 54, 56, 58
 - starting the navigation process, 55, 59
 - types and their uses, 97
 - variables, 65, 65
- master map (auto_master)
 - /- mount point, 53, 57
 - comments in, 54
 - comparing with /etc/mnttab file, 52
 - contents, 53, 55
 - description, 98
 - overview, 53, 54
 - preinstalled, 100
 - security restrictions, 106
 - syntax, 54
 - when to run automount command, 98
- mirror mounts
 - mounting all file systems from one server, 77
 - overview, 49
- mnttab file
 - comparing with auto_master map, 52
- modifying
 - NFS referrals, 91
- mount command, 114
 - autofs and, 14
 - diskless clients' need for, 14
 - failover with, 118, 118
 - manually mounting file systems, 76
 - NFS URL, 80
 - NFS URL with, 118
 - options
 - description, 115
 - no arguments, 119
 - public, 79
 - using, 117
- mount of *server:pathname* error, 149
- mount points

- /- as master map mount point, 53, 57
 - /home, 53, 55
 - /net, 55
 - /nfs4, 53, 55
 - avoiding conflicts, 99
 - mountall command, 121
 - mountd daemon, 163
 - checking response on server, 143
 - not registered with rpcbind, 152
 - verifying if running, 144, 152
 - mounting
 - all file systems in a table, 121
 - autofs and, 14, 61
 - background retries, 115
 - diskless client requirements, 14
 - examples, 117
 - FedFS, 91
 - force direct I/O, 115
 - foreground retries, 115
 - keyboard interruption during, 133
 - mirror mounts and, 49
 - nfsd daemon and, 39
 - overlying already mounted file system, 118
 - portmapper and, 39
 - public file handle and, 40
 - read-only specification, 117, 117
 - read-write specification, 117
 - remote mounting
 - daemons required, 133
 - troubleshooting, 141, 144
 - soft compared to hard, 133
 - mounting file systems
 - autofs and, 77
 - boot time method, 75
 - disabling access for one client, 78
 - manually (on the fly), 76
 - mounting all from one server, 77
 - NFS URL with, 79
 - overview, 74
 - task map, 74
 - through a firewall, 79
- N**
- /net mount point, 55
 - /nfs4 mount point, 53, 55
 - namespaces
 - accessing shared, 104
 - autofs and, 22
 - naming services
 - autofs map maintenance methods, 98
 - navigating using maps
 - overview, 53
 - starting the process, 55, 59
 - negotiation
 - file transfer size, 39
 - WebNFS security, 21
 - netconfig file
 - description, 158
 - Network Lock Manager, 19
 - NFS
 - commands, 111
 - daemons, 160
 - version negotiation, 28
 - NFS ACL
 - description, 18, 37
 - error message, Permission denied, 154
 - NFS administration
 - administrator responsibilities, 71, 95
 - NFS can't support nolargefiles message, 153
 - NFS clients
 - incompatible operating system support, 105
 - NFS services, 15
 - NFS environment
 - Secure NFS system, 45
 - NFS locking
 - client-side failover and, 42
 - NFS referrals
 - creating, 90, 94
 - overview, 50
 - removing, 91
 - NFS server logging
 - enabling, 73
 - overview, 22
 - NFS servers
 - autofs selection of files, 64

- daemons required for remote mounting, 133
- identifying current, 145
- maintaining, 71, 95
- replicating shared files, 106
- troubleshooting
 - clearing problems, 141
 - remote mounting problems, 140, 153
- weighting in maps, 65
- NFS services
 - restarting, 145
 - selecting different versions on client by
 - changing the SMF properties, 83
 - using the mount command, 84
 - selecting different versions on server, 82
 - task map, 81
- NFS troubleshooting
 - determining where NFS service has failed, 144
 - hung programs, 154
 - remote mounting problems, 153
 - server problems, 141
 - strategies, 133
- NFS URL
 - autofs and, 107
 - mount command example, 118
 - mounting file systems with, 79
 - mounting with, 21
 - syntax, 89
 - WebNFS and, 88
- NFS V2 can't support largefiles message, 153
- NFS Version 4
 - features in, 28
- nfs4cbd daemon, 164
- nfscast: cannot receive reply message, 149
- nfscast: cannot send packet message, 149
- nfscast: select message, 150
- nfsd daemon, 164
 - checking response on server, 142
 - mounting and, 39
 - verifying if running, 144
- nfslog.conf file, 159
- nfslogd daemon
 - description, 165
- nfslogd file, 158
- nfsmapid daemon
 - ACLs and, 37
 - additional information about, 172
 - configuration files and, 167
 - configuring the NFSv4 default domain, 170
 - description, 17, 165
 - DNS TXT records and, 168
 - identifying NFSv4 domain, 169
 - precedence rules and, 167
- NFSMAPID_DOMAIN keyword, 38
- nfsmapid_domain parameter, 166
- nfsref command
 - description, 132
 - example, 94
- nfsstat command, 134, 145
- NIS name service
 - updating autofs maps, 98
- no info message, 150
- No such file or directory message, 152
- nobrowse option
 - auto_master file, 108
- nobrowse parameter
 - setting, 108
- nolargefiles option
 - error message, 153
 - mount command, 115
- nosuid option
 - share command, 126
- Not a directory message, 149
- Not found message, 147
- nsdb-list command
 - description, 132
- nsdb-nces command
 - description, 132
- nsdb-resolve-fsn command
 - description, 132
- nsdb-update-nci command
 - description, 132
 - example, 92
- nsdbparams command
 - description, 132
 - example, 93

- nthreads option
 - lockd daemon, 163
- number sign (#)
 - comments in direct maps, 56
 - comments in indirect maps, 58
 - comments in master map (auto_master), 54
- O**
- O option
 - mount command, 118
- o option
 - mount command, 117
 - share command, 125, 128
- open errors
 - NFS and, 16
- OPEN share support
 - NFS Version 4, 34
- operating systems
 - map variables, 65
 - supporting incompatible versions, 105
- OSNAME map variable, 65
- OSREL map variable, 65
- OSVERS map variable, 65
- overlying already mounted file system, 118
- P**
- passwords
 - autofs and superuser passwords, 14
 - DH password protection, 46
- pathconf: no info message, 150
- pathconf: *server not responding* message, 150
- Permission denied message, 153
- permissions
 - NFS Version 3 improvement, 16
- plus sign (+)
 - in autofs map names, 66, 67
- portmapper
 - mounting and, 39
- pound sign (#)
 - comments in direct maps, 56
 - comments in indirect maps, 58
 - comments in master map (auto_master), 54
- printing
 - list of remotely mounted directories, 131
 - list of shared or exported files, 131
- problems with ACLs in NFS
 - avoiding, 38
- processor type map variable, 65
- programs
 - hung, 154
- projects
 - consolidating files, 102
- pstack command, 137
- public file handle
 - autofs and, 107
 - mounting and, 40
 - NFS mounting with, 21
 - WebNFS and, 88
- public option
 - in dfstab file, 89
 - mount command, 79, 116
 - share error message, 155
 - WebNFS and, 88
- public-key cryptography
 - common key, 47
 - conversation key, 47
 - database of public keys, 46, 47
 - DH authentication, 47, 47
 - secret key
 - database, 47
 - deleting from remote server, 48
 - time synchronization, 47
- public-key map
 - DH authentication, 47
- R**
- r option
 - mount command, 117
 - umountall command, 121
- read-only type
 - file selection by autofs, 61, 64
 - mounting file systems as, 117, 117

- sharing file systems as, 125, 125, 128
- read-write type
 - mounting file systems as, 117
 - sharing file systems as, 125, 128
- referrals *See* NFS referrals
- remote file systems
 - listing clients with remotely mounted file systems, 131
 - unmounting groups, 121
- remote mounting
 - daemons required, 133
 - troubleshooting, 140, 144
- Remote Procedure Call (RPC)
 - Secure
 - overview, 46
- remount message, 148
- removing
 - NFS referrals, 51, 91
- removing locks, 113
- reparsed daemon, 172
- replicas must have the same version message, 155
- replicated file system, 42
- replicated mounts
 - soft option and, 155
- replicated mounts must be read-only message, 155
- replicated mounts must not be soft message, 155
- replicating shared files across several servers, 106
- restricting
 - displayed file system information, 80
- rlogin command
 - secure NFS and, 48
- ro option
 - mount command, 117
 - mount command with `-o` flag, 117
 - share command, 125, 128
- root directory
 - mounting by diskless clients, 14
- root option
 - share command, 127
- RPC

- authentication, 47
 - Secure
 - DH authorization issues, 48, 48
- rpcbind daemon
 - dead or hung, 152
 - mountd daemon not registered, 152
- rpcinfo command, 137
- RPCSEC_GSS, 21
- rw option
 - mount command, 117
 - share command, 125, 128
- rw=client option
 - umountall command, 125

S

- s option
 - umountall command, 121
- secret key
 - database, 47
 - deleting from remote server, 48
 - server crash and, 48, 48
- Secure NFS system
 - administering, 85
 - DH authentication and, 85
 - domain name, 85
 - overview, 45
- Secure RPC
 - DH authorization issues, 48, 48
 - overview, 46
- security
 - applying autofs restrictions, 107
 - DH authentication
 - overview, 47, 47
 - password protection, 46
 - user authentication, 45
 - file-sharing issues, 125, 127
 - NFS Version 3 and, 16
 - Secure NFS system
 - administering, 85
 - overview, 45
 - Secure RPC
 - DH authorization issues, 48, 48

- overview, 46
 - UNIX authentication, 45, 47
 - security and NFS
 - description, 18, 37
 - error message, *Permission denied*, 154
 - security flavors, 21
 - security mode selection and `mount` command, 117
 - serial unmounting, 121
 - server not responding message, 148, 150
 - hung programs, 154
 - keyboard interrupt for, 133
 - remote mounting problems, 152
 - `server_delegation` parameter, 165
 - `server_versmax` parameter, 165
 - `server_versmin` parameter, 164
 - servers, 106
 - See also* NFS servers
 - autofs selection of files, 61
 - crashes and secret keys, 48, 48
 - home directory server setup, 101
 - NFS servers and `vfstab` file, 75
 - NFS services, 15
 - servers and clients
 - NFS service, 15
 - `setfacl` command
 - NFS and, 37
 - setgid mode
 - `share` command, 126
 - setting
 - `nobrowse` parameter, 108
 - setuid mode
 - Secure RPC and, 48
 - `share` command, 126
 - `share` command
 - description, 124
 - enabling WebNFS service, 89
 - options, 125
 - security issues, 127
 - `shareall` command, 130
 - sharing *See* file sharing
 - `showmount` command, 131
 - example, 80
 - `showmount_info` property, 80
 - single-user mode and security, 48
 - slash (/)
 - /- as master map mount point, 53, 57
 - master map names preceded by, 54
 - root directory, mounting by diskless clients, 14
 - `snoop` command, 139
 - soft option
 - `mount` command, 117
 - special characters in maps
 - enclosing in quotation marks, 70
 - `statd` daemon, 172
 - superuser
 - autofs and passwords, 14
 - synchronizing time, 47
- T**
- t option
 - `lockd` daemon, 162
 - TCP
 - NFS Version 3 and, 19
 - `telnet` command
 - secure NFS and, 48
 - time synchronization, 47
 - transport protocol
 - NFS negotiation, 38
 - transport setup problem
 - error message, 151
 - Troubleshooting
 - NFS Commands, 174
 - troubleshooting
 - autofs, 146
 - avoiding mount point conflicts, 99
 - error messages generated by automount - v, 146
 - miscellaneous error messages, 148
 - NFS
 - determining where NFS service has failed, 144
 - hung programs, 154
 - remote mounting problems, 140, 153
 - server problems, 141
 - strategies, 133

truss command, 140

U

/usr directory

mounting by diskless clients, 14

/usr/kvm directory

mounting by diskless clients, 14

/usr/lib/fs/nfs/fedfs-11.schema file, 92

/usr/sbin/mount command *See* mount command

/usr/sbin/nsdb-list command

description, 132

/usr/sbin/nsdb-nces command

description, 132

/usr/sbin/nsdb-resolve-fsn command

description, 132

/usr/sbin/nsdb-update-nci command

description, 132

/usr/sbin/nsdbparams command

description, 132

/usr/sbin/showmount command, 131

/usr/sbin/unshareall command, 130

UDP

NFS and, 19

umount command

autofs and, 14

description, 120

umountall command, 121

UNIX authentication, 45, 47

unmapped user or group IDs

checking for, 38

unmounting

autofs and, 14, 61

examples, 120

groups of file systems, 121

mirror mounts and, 49

unshare command, 129

unshareall command, 130

unsharing and resharing

NFS Version 4, 29

unsharing file systems

unshare command, 129

unshareall command, 130

URL service types

WebNFS and, 90

V

-V option

umount command, 120

-v option

automount command, 146

variables in map entries, 65, 65

verifiers

RPC authentication system, 46

version negotiation

NFS, 28

vfstab file

automount command and, 53

enabling client-side failover, 78

mounting by diskless clients, 14

mounting file systems at boot time, 75

NFS servers and, 75

volatile file handles

NFS Version 4, 32

W

WARNING: *mountpoint* already mounted on message, 148

WebNFS service

browsing, 89

description, 43

enabling, 88

firewalls and, 90

overview, 20

planning for, 87

security negotiations and, 21

task map, 87

URL service types and, 90

weighting of servers in maps, 65

write errors

NFS and, 16