

---

**CONVEX**  
**Network File System Utilities**  
**V10.0**  
**Release Notice**

---



Document No. 710-003730-009

December 1991

---

**CONVEX**  
**Network File System Utilities V10.0**  
**Release Notice**

Document No. 710-003730-009

Copyright 1991 CONVEX Computer Corporation.  
All rights reserved.

This document is copyrighted. This document may not, in whole or part, be copied, duplicated, reproduced, translated, electronically stored, or reduced to machine readable form without prior written consent from CONVEX Computer Corporation.

Although the material contained herein has been carefully reviewed, CONVEX Computer Corporation (CONVEX) does not warrant it to be free of errors or omissions. CONVEX reserves the right to make corrections, updates, revisions or changes to the information contained herein. CONVEX does not warrant the material described herein to be free of patent infringement.

Unless provided otherwise in writing with CONVEX Computer Corporation (CONVEX), the program described herein is provided as is without warranty of any kind, either expressed or implied, including, but not limited to the implied warranties of merchantability and fitness for a particular purpose. Some states do not allow the exclusion of implied warranties. The above exclusion may not be applicable to all purchasers because warranty rights can vary from state to state. In no event will CONVEX be liable to anyone for special, collateral, incidental or consequential damages, including any lost profits or lost savings, arising out of the use or inability to use this program. CONVEX will not be liable even if it has been notified of the possibility of such damage by the purchaser or any third party.

CONVEX and the CONVEX logo ("C") are registered trademarks of CONVEX Computer Corporation.

ConvexOS is a trademark of CONVEX Computer Corporation.

NFS and SunOS are trademarks of Sun Microsystems, Inc.

Ethernet is a trademark of Xerox Corporation.

Printed in the United States of America

---

## Introduction

This document describes the V10.0 Release of the CONVEX Network File System (NFS) software. Always refer to this release notice before reporting problems: your questions may be answered herein. Fixes and workarounds listed in this document may help you avoid rediscovering known problems.

The CONVEX NFS release includes the following functionality:

- NFS (Network File System) – provides transparent access to remote file systems in a heterogeneous network.
- NIS (Yellow Pages) – provides a distributed network lookup service.
- NETdisk – allows diskless workstations to boot from a CONVEX system.
- RPC/XDR (Remote Procedure Call/External Data Representation) – a set of libraries which implement network transactions. The NFS products are built on the RPC/XDR libraries.
- RPC Protocol Compiler (rpcgen) – a tool which generates C code to implement the RPC protocol.
- REX (Remote EXecution) – a program which executes a command remotely using NFS.
- Automount - a program that allows on-demand mounting of NFS file systems.
- Secure NFS - a version of NFS whose transaction requests are protected by an encryption protocol.

---

### → Note

---

**The Secure NFS functionality is not available to international customers outside of the U.S. and Canada at this time due to legal restrictions on the export of encryption algorithms by the U.S. Department of State's Office of Munition Control.**

The remaining sections in this document include the following information:

- the contents of this distribution.
- notes and cautions about use of the software.
- known restrictions in the software.
- bugs fixed in this release.
- known software problems.
- the release documentation.

To install the software described in this document, consult the *Installation Procedures, ConvexOS and Utilities V10.0*.

---

## Contents of This Distribution

The distribution package for this release of CONVEX Network File System Utilities consists of

- This release notice
- Distribution media for the software
- Documentation

The specific contents of the software and documentation distribution are described below.

---

### Software for USA Distribution

Item	Qty	Type	Part Number	Description
1.	1	Mag Tape	710-009615-002	ConvexOS Utilities, USA V10.0

---

### Software for International Distribution

Item	Qty	Type	Part Number	Description
1.	1	Mag Tape	710-009515-002	ConvexOS Utilities, INTL V10.0

---

## Notes and Cautions

This section contains useful information and words of caution about the product.

- The V10.0 release of the CONVEX Network File System is bundled on the same tape as the V10.0 release of ConvexOS and Utilities. To extract this product, you must know the activation key for your system. If you have not received an activation key with the product, contact the CONVEX Technical Assistance Center for instructions. In the continental United States, call 1 (800) 952-0379.

Prerequisites for this release of the CONVEX Network File System Utilities are as follows:

- your system must be running the V10.0 release of the ConvexOS operating system
  - your system must be running the V10.0 release of the ConvexOS Utilities
  - your system must be running the V10.0 release of the CONVEX Internet Services
- If you received the USA distribution tape, Secure NFS will automatically be installed when you select the CONVEX Network File System Utilities for installation.

To use NFS on your system, start up system daemons in the `/etc/rc.local` system start-up file, copying lines from `/etc/rc.local.v10.0`. See the *CONVEX Network File System System Manager's Guide* for complete instructions on how to set up NFS, the automounter, Secure NFS and NIS for your site.

---

## New Features

The lock manager has been upgraded to the NFS 4.1 version. With this version of the lock manager, all local file locking is done within the Operating system. The lock manager daemon (`rpc.lockd`) is now only involved in locking files accessed via `nfs` and serving lock requests from remote systems.

---

## Interoperability with non-Convex NIS (yp) masters

If you are using a non-CONVEX machine as the master NIS server (such as a machine from Sun Microsystems), change the NIS configuration on that machine. Convex NIS expects the NIS services map and `rpc` map to be built so that it may be searched quickly with a single `yp_match` access. Sun and some other vendors do not build these maps, and the master NIS Makefile must change to build the map so as to allow the Convex to find the information it needs. This change will not interfere with NIS accesses by machines from other vendors, as it only supplies additional information for the CONVEX.

To make this change, copy lines from the file `/usr/etc/yp/Makefile` on the CONVEX. Find the string `services.time`, copy all lines down to the string `protocols.time` to a temporary file and move the temporary file onto the non-CONVEX machine. Add the contents of the temporary file to the master YP Makefile on the non-CONVEX machine.

On a Sun running SunOS 4.0 or greater, look for the Makefile in `/var/yp`, while on older Suns look for it in `/etc/yp`. Go to that directory and replace the appropriate lines in the Makefile with those from the temporary file. Once the Makefile is ready, remove the files `services.time` and `rpc.time`. Build the new maps with `make`. The CONVEX can now access the services and `rpc` maps correctly.

Do not set the system domainname if you do not plan on running the NIS system. In general, if `ypbind` is running, the domainname must be set; otherwise the domainname should not be set.

---

## Interoperability between NIS (yp) and sendmail

The default sendmail configuration does not support the use of the YP/NIS alias map. The result of this is that sending mail to a YP alias will fail unless the sendmail configuration file is altered.

The steps required to enable YP aliases within sendmail are:

1. Search in the file `/usr/lib/sendmail.cf` for a line that begins with the characters "OA". An example would be:  
`OA/usr/lib/aliases`
2. Add the line "Op" immediately after the OA line. An example would look like:  
`OA/usr/lib/aliases`  
`Op`
3. Refreeze the sendmail configuration file by executing a `"/usr/lib/sendmail -bz"`.

After performing step three, YP mail aliases should be enabled.

---

## Altering NIS (yp) netgroup maps and mountd

If the NIS(yp) netgroup map on a running system is altered, the mount daemon (mountd) will not see the changes until the mount daemon is killed and restarted or the system is rebooted.

The procedure to kill and restart the mountd is as follows:

1. find the pid of the mountd "ps -aux | grep mountd"
2. kill -1 <the pid found in step one>.
3. find the pid of the Internet services daemon "ps -aux | grep inetd"
4. kill -1 <the pid found in step one>.

All mounts made after this point will have knowledge of the changes to the netgroup YP/NIS map.

## Interoperability between NIS and the Name Server

The NIS server normally uses the YP maps to answer all queries, which can be a problem in networks using the Berkeley Internet Name Domain Server (BIND). If desired, NIS can be configured to use BIND to resolve host information. This allows simpler configuration of NIS client machines, since these machines can get accurate hosts information without being configured to use BIND themselves. This procedure works on both ConvexOS V9.1 and V10.0. See Chapter 9 of the *Internet Services System Manager's Guide* for more information on BIND.

To configure NIS to use BIND, these changes are necessary: all NIS server daemons must be restarted with a new command line option; the file `/usr/etc/yp/Makefile` must be changed to build the NIS `hosts.byname` and `hosts.byaddr` maps differently; and the `hosts.byname` and `hosts.byaddr` maps must be rebuilt and pushed out to all NIS servers.

First, edit `/etc/rc.local` on all NIS server machines and add `-i` to the invocation of `/usr/etc/ypserv` as follows:

```
if [ -f /usr/etc/ypserv -a -d /etc/yp/`bin/domainname` ]; then
    /usr/etc/ypserv -i & echo -n ' ypserv'
fi
```

The `-i` flag ensures that `ypserv` will attempt to use BIND without regard to the `/etc/use_nameserver` file. Next, edit `/usr/etc/yp/Makefile` to add `-b` to the invocation of `makedbm` in the lines shown below:

```
hosts.time: $(DIR)/hosts
    -@if [ -f $(DIR)/hosts ]; then \
        umask 022; \
        sed -e "/^#/d" -e s/#.*$$// $(DIR)/hosts | $(STD-
HOSTS) \
        | awk '{for (i = 2; i <= NF; i++) print $$i, $$0}' \
        | $(MAKEDBM) -b - $(DOMDIR)/hosts.byname; \
        $(STDHOSTS) $(DIR)/hosts | \
        awk 'BEGIN {OFS="\t";} $$1 !~ /^#/ {print $$1,
$$0}' \
        | $(MAKEDBM) -b - $(DOMDIR)/hosts.byaddr; \
        touch hosts.time; \
        echo "updated hosts"; \
```

This change adds a special `YP_INTERDOMAIN` key to the NIS hosts maps to tell the YP server that it should make host database lookups via BIND. Finally, use the following steps to update and push out the new NIS maps:

```
# cd /usr/etc/yp
# rm hosts.time
# make
```

---

### → Warning

---

Do not distribute these maps to CONVEX NIS servers which are not set up to use BIND by default (i.e., which do not have the file `/etc/use_nameserver`) until they are running `ypserv` with the `-i` flag. Otherwise, the YP server will enter an endless loop.

---

## NETdisk changes

NETdisk has undergone some major improvements to permit it to work with recent SunOS release tapes. The use of the NETdisk tools has not changed, though users should be aware of the following:

- The support for sharing of non-executable files such as 'man' pages is no longer supported by these scripts since Sun Microsystems has dropped their support for this feature. If you were using this feature in the past, you may encounter difficulties if you try to add new clients.
- Though Sun native installs now support multiple versions and multiple architectures, configurations such as these are not supported by Convex. Such installations fundamentally work, but the root prototype tree installed in /usr/etc/install/proto is shared for all architectures, which may cause problems with multiple versions in particular.
- The SunOS 4.1.1 release is the first Sun release known to contain compressed 'tar' archives. Convex does not supply 'uncompress', the tool needed to work with these archives, at this time. The 'compress' tools are in the public domain and available from many sources, including the Convex Users Group tape. If you run INSTALL(8) with 'uncompress' in your search path, you will be able to install SunOS 4.1.1 successfully.

Convex has tested these scripts to work correctly with the following SunOS release tapes:

- |               |               |             |           |
|---------------|---------------|-------------|-----------|
| • SunOS 4.0   | Sun-3(68020)  | 700-1577-10 | cartridge |
| • SunOS 4.0.3 | Sun-3(68020)  | 700-2159-10 | 9-track   |
| • SunOS 4.1   | Sun-3(68020)  | 700-1877-14 | cartridge |
| • SunOS 4.1   | Sun-4c(SPARC) | 700-2521-10 | cartridge |
| • SunOS 4.1.1 | Sun-3(68020)  | 790-4435-02 | 9-track   |

---

## Documentation

Current documentation for CONVEX NFS Utilities V10.0 includes the following manuals and addenda:

- *CONVEX NFS Concepts*, First Edition
- *CONVEX NFS System Manager's Guide*, Third Edition
- *CONVEX NFS Reference Set*, Third Edition.
- *CONVEX NFS Programmer's Reference*, Second Edition.

---

## NFS Module Names

This section describes the module names used in the following sections describing restrictions, bugs and bug fixes in CONVEX NFS.

Module Name	Short Description
rpc	User and kernel level rpc and xdr
srpc	User and kernel level secure RPC
auto	Automounter, automount(8)
lckmgr	Network based lock manager
yp	Yellow Pages
rpcgen	RPC protocol compiler, rpcgen(1)
rquota	Remote file system quotas, rquotad(8c)
nfs	Network File System
rex	Remote Execution, rex(1c) and rexd(8c)
rwall	Remote write to all users
rusers	rusers(1c), rusersd(1c)
rstat	Remote system status
netdisk	NETdisk
netmgt	Configuration and Management
man	CONVEX NFS man pages
cncpts	CONVEX NFS Concepts manual
refer	CONVEX NFS Reference set
smg	CONVEX NFS System Manager's Guide
index	CONVEX NFS Master Index

This section describes the bug fixes that have been made since the last release of CONVEX NFS.

---

### auto

PR-17914

Running automount without NIS configured results in the following error messages being logged:

```
<date hostname> automount[10346]: YP bind failed: args to yp function are bad  
<date hostname> automount[10346]: auto.master: args to yp function are bad
```

Resolution: The automounter has been changed to check if NIS (yp) is running before calling any NIS functions.

PR-17313 PR-20108

The automounter can hang with the following deadlock. The automounter is in `clntkudp_callit_addr()` doing a `getattr` on a remote file system. The function `clntkudp_callit_addr()` is waiting to get access to the output buffer for the client handle being used for the request. However, the output buffer is referenced by an `mbuf` which is on the automounter's socket receive buffer.

Resolution: The problem has been fixed in `clntkudp_callit_addr()` by waiting on the client handle output buffer to become free before returning from `clntkudp_callit_addr()`. Note that this is the method used by Sun NFS.

PR-20166 PR-20434

When attempting to mount exported filesystems from the NOS/VE system via the automounter the mounts fail. Normal mounts work fine.

Resolution: The automounter has been modified to use a different method of pinging the remote mountd.

PR-20369

If the mount point specified in an automounter direct map contains a symlink, the mount point will not be resolved to the real `/tmp_mnt...` mount point when it is accessed. The mount point appears to be a link to nowhere and users can't access the files in the server directory via the mount point path.

Resolution: The automounter has been modified to allow the map entry paths (mount points) to contain symlinks.

PR-21466

The man page for `mount` states that the `grpuid` option can be used for setting the group ID for a file the same as the group ID of the directory (default for the CONVEX). This does not work for a mounted file system on a Sun (files reside on the Sun). Setting the `SETGID` bit on the Sun directory gives the correct result.

Resolution: The automounter has been modified to make the `'grpuid'` option operate

correctly.

PR-22310

The automounter leaves a link to nowhere: lrwxrwxrwx xxxxxx xxxxxx-> when the directory it is trying to mount is not exported to the machine on which the automounter is running.

Resolution: Error logging was added to the automounter so that the case of trying to mount a remote directory which is not exported or which does not contain the client in its access list, will be easier to diagnose. The empty symbolic link for the mount-point is the correct behavior and is the way the SunOS 4.1 automounter handles the same situation. The error logging is displayed if the verbose command line parameter is supplied when starting the automounter.

PR-22980

The automounter will core dump if an indirect map of 100+ entries is used.

Resolution: The automounter has been fixed so that it does not corrupt its process stack when processing a large indirect map.

---

## install

PR-20023

When NFS is selected for optional product installation from an International Utilities tape, the install script checks for disk space needed for Secure NFS on the SPU, and (if applicable) complains: \*\*\* Insufficient space on the spu for Secure NFS Please free 2Mb on /mnt/os. Note that Secure NFS is not even on the international tape, contrary to what the install notes suggest. At least there is no way to select Secure NFS.

Resolution: The install scripts will no longer check available SPU space unless a secure NFS sysgen is going to be performed during the product installation.

---

## lckmgr

PR-14224

The lock daemon's total SZ and RSS seem to grow proportionately to the total time used by the process. This seems indicative of a memory leak due to an improperly free()ed malloc().

Resolution: The lock manager has been greatly changed for ConvexOS V10.0 and the known memory leaks are now fixed.

PR-15619

The `fcntl(2)` system call can be interrupted when stopped and started while blocked, waiting for a file lock resulting in the failure of the lock request.

Resolution: Blocking `fcntl()` file lock requests can now be stopped and started without interrupting the `fcntl()` system call.

PR-19385 PR-20219 PR-20967

The `rpc.lockd` daemon aborted and dumped core.

Resolution: The lock manager has been altered and restructured for ConvexOS V10.0, resulting in a much more stable lock manager.

PR-21114

If a file resides on machine A and a process that runs on machine B (which NFS-mounts the filesystem containing the file) locks the file with `lockf(fd, F_TLOCK, 180)`, and machine B crashes, it is not possible to lock the file from either system after B is rebooted because of permission problems.

Resolution: This problem has been fixed.

---

## man

PR-19329

The "SEE ALSO" section of the `publickey.3r` man page has a line truncated to: "RPC Programmer's Manual in"

Resolution: The man page for `publickey(3r)` has been corrected to refer to the RPC Programming Guide in the Convex NFS Reference Set.

PR-20041

The `rpc.statd` program supports the `-dlevel` option, but the man page makes no mention of any options. The `-dlevel` option needs to be explained on the man page.

Resolution: The `statd.8c` man page has been updated to include descriptions of the `-d` and `-D` command line parameters.

PR-21572

The man pages for `domainname(8)`, `getdomainname(2)` and `setdomainname(2)` do not mention that they are for NIS(YP) domains as opposed to the Internet Domain Name. The NAME and DESCRIPTION sections should have NIS added before the word "domain".

Resolution: The man pages `domainname.1` and `getdomainname.2` have been updated to clarify that the domainname referred to is the yellow pages domain.

PR-18005

The intro.3r man page refers to non-existent man pages. The following are referred to but have no man pages: klm\_prot (3r)nlm\_prot (3r) sm\_inter.3r yp (3r)

Resolution: Added the man pages for klm\_prot.3r, nlm\_prot.3r, sm\_inter.3r, yp.3r and edited the descriptions of the standard RPC services in intro.3r.

PR-18361

The exports(5) and netgroup(5) man pages should state that to use netgroups you must also be running and using NIS(YP).

Resolution: A note has been added to the man pages exports.5 and netgroups.5 indicating that netgroup relies on NIS(YP).

---

## netdisk

PR-19890

The netdisk install script "setup\_client" does not work unless the system is running NIS (Yellow Pages).

Resolution: setup\_client has been modified to get the domain name by executing /bin/domainname instead of parsing the rc.local file. If the YP service is to be used during the client installation, the domain name must be set and yp must be running when setup\_client is executed. If the domain name is not set, YP cannot be running and will not be used by setup\_client.

If there is a domain name configured, setup\_client will use 'ypmatch' and thus, the domain name \_must\_ be set and yp \_must\_ be running when setup\_client is executed.

---

## netmgt

PR-17905

The permissions on /etc/rmtab are not secure enough. It seems the permissions should be rw-r--r-- but not rw-rw-rw.

Resolution: The permissions of /etc/rmtab have been changed to be rw-r--r--.

PR-19011

The /etc/rc.\* files are currently setup to perform NFS mounts before they have started their own NFS server daemons(mountd). The portmap (mountd) daemon should be started before any NFS client requests are performed. This would alleviate the problem with systems that are both NFS servers and clients, that must timeout the NFS mounts before they can start the NFS server daemons.

Resolution: The start-up of inetd has been moved to rc.local from rc.std. Now, inetd will be running and able to start mountd before NFS mounts are attempted.

PR-19420

When a filesystem is exported with a `-anon=-1`, the authentication fails during the mount.

Resolution: The authentication now functions properly with `-anon=-1`.

---

## **nfs**

PR-18795

Inconsistencies in the treatment of uids in the Unix kernel result in problems with NFS access when a negative uid is used. In particular, a directory can be created but the permissions will not allow use of that directory. This is a problem with root access when the filesystem is exported with `"anon=-2"` and with access by PC clients, which often use uid `-2`.

Resolution: File access with a negative uid is now consistent in its behavior.

PR-18798

An NFS operation is leaving files named `.nfs??????` around which can take up a large amount of disk space. It appears to be caused by programs that unlink temporary files.

Resolution: NFS operations on open and unlinked files now properly clean up upon completion.

PR-18502

If the `udpcksum` tunable parameter is set to 1 in the `bootcmd.local` file, NFS will no longer work.

Resolution: There were two problems: 1. setting wrong IP source add before UDP checkless calculation, 2. overwriting IP header length when calculating UDP checkless. Both problems are fixed.

PR-18842

Unlinking open files over NFS on a machine running bids can in some cases result in a system hang.

Resolution: The problem was fixed by not inserting unlinked/open files into the directory name lookup cache.

PR-20011

When a `"mv"` is performed of a file as large or larger than the amount of physical memory in the system to an NFS mounted fleshiest, all system memory can potentially be allocated to buffering data for the `"me"` command and the system can deadlock.

Resolution: NFS operations are no longer capable of buffering an unlimited amount of memory during transfer operations.

PR-20460

The NFS biod daemons can potentially deadlock during interaction with the virtual memory system.

Resolution: The deadlock has been eliminated.

PR-20566

Killing an nfs process can result in a system hang if it is killed while the NFS server is executing client requests.

Resolution: The nfs process now shuts down properly when it is terminated.

PR-20639

An nfsd process can hang during a read directory operation if the directory entry on the filesystem is corrupted.

Resolution: The nfs code that reads directories now does sanity checks on some aspects of the directory information passed to it from the filesystem.

PR-21692

The owner of a file which does not have read or write permission for the owner cannot read or write the file over NFS if the owner is -2 (for anon=-2 mounts).

Resolution: The problem has been fixed.

PR-22542

The NFS server code defines a global linked list (rfsfreesp) which is not protected by a semaphore. This problem could potentially cause corruption of the rfsfreesp list and a system panic.

Resolution: The list rfsfreesp is now protected by a semaphore.

PR-22585

A table involved in nfs/rpc authentication became corrupted. The corruption caused a system panic.

Resolution: The source of the authentication table corruption has been fixed.

PR-22822

There is currently a limit on how long an NFS access list can be in the /etc/exports file. If a list is specified that exceeds the limit, no error is returned and the behavior is the same as if there were no access list.

Resolution: The exportfs command has been modified to display a warning if an entry in the /etc/exports file is too long.

PR-21347

The NFS server code does not support read-dir request sizes of less than 512 bytes.

Resolution: Support for read-dir requests less than 512 bytes has been added to the NFS server code.

---

## **passwd**

PR-19919

If 'root' runs passwd to change the root password, the passwd application wants to use the new password to encrypt the secret key for root in the yp map publickey.byname. The string variable 'newpass' is passed to the reencrypt\_secret() function. But in the main() function, the newpass string variable is not initialized if the user id is 0. Thus, the secret key for root will be encrypted using a non-valid password and is then stored in the yp map.

Resolution: /bin/passwd was changed to properly initialize the string used to re-encrypt root's secret DES key in the publickey.byname YP map.

---

## **rpcgen**

PR-22139

The rpcgen utility should check for read/write failures and call perror() if an error occurs.

Resolution: Error checking and diagnostic messages have been added to rpcgen in the case of failed reads and writes.

---

## **rpc**

PR-21877

A User process can hang on the Client side when accessing a file over NFS over the Ultranet uip interface.

Resolution: The problem has been fixed.

---

## **srpc**

PR-20466

Files including /usr/include/des\_crypt.h will not compile correctly.

Resolution: The file /usr/include/des\_crypt.h has been fixed so that files including it will compile correctly.

---

## yp

PR-10828 PR-11788

ypserv and ypbind should use syslog to report error conditions.

Resolution: The ypbind and ypserv daemons now report all errors via syslog messages.

PR-19918

When an ordinary user changes his or her password with yppasswd, yppasswd wants to encrypt the users secret key in the publickey.byname yp map using the new password. yppasswd does a getnetname() to get the name of the user doing the password change. The library function getnetname() uses the process' effective user id to determine which user name to return. The yppasswd application is running with the effective user ID of root since the file mode is suid as root.

Resolution: The correct effective user ID is now set by yppasswd before getnetname() is called to get the users network name.

PR-20674

The yppasswdd daemon does not execute a mkpasswd after updating the /etc/passwd file, thus local logins and applications that use the password dbm database files fail.

Resolution: The yp password daemon (rpc.yppasswdd) has been modified to add the 'mkpasswd' command to the command string which does the ypmake if the -m flag is supplied.

PR-22512

The system accounting utility /etc/sa will dump core when executed on an NIS(yp) client.

Resolution: The problem has been fixed.

---

## Known Software Problems

This section lists the known problems with the CONVEX Network File System Utilities software as of December, 1991. Problems reported after this date may not be reflected in the document. Please refer to this list prior to reporting a problem to ensure that it has not been previously reported. Serious problems include workarounds if they are known.

---

### netdisk

PR-18535 PR-18535

If the SunOS installation is done in a fresh, cleaned-out file system and /tftpboot points to /filesystem/tftpboot but the directory isn't there, the installation of the client fails.

Workaround: Do mkdir /filesystem/tftpboot and retry the installation.

PR-13180

The daemon rpc.bootparamd can dump core if it fails in an attempt to read kernel memory.

PR-23090

If the search path is sufficiently strange, programs called from INSTALL can fail to find binaries they need. The reason is that INSTALL is not exporting its PATH variable.

---

### netmgf

PR-11036

The command 'rpcinfo -u host nlockmgr' produces the following output: program 100021 version 1 ready and waiting rpcinfo: RPC: Procedure unavailable program 100021 version 2 is not available

PR-20399

Mountd should be started automatically by inetd if it has aborted.

---

### nfs

PR-08924

Reading and writing to a named pipe can hang forever if the named pipe is on different file systems but is the same underlying object. Somehow, because of NFS, the system does not realize it is the same object. However, NFS is suppose to be transparent to the user.

PR-14729

If the "noac" flag is used to do an NFS mount from a X9.0 system to a V8.1 system, then a "trace/BPT trap" message will appear when an executable is run.

PR-15080

When a filesystem is exported to a specific list of machines, many of which are HP workstations, it is exported -rw, and the HPs mount it -rw. The first 16 machines in the host list work correctly, but all HP workstations after the 16th machine do not have write access to the filesystem. Other CONVEX machines get write access to the filesystem regardless of its position in the host list.

PR-15585 PR-19430

The library version of chkpnt is having problems with fdpath failing under certain circumstances. This appears to be related to NFS mounting the filesystems.

PR-19297 PR-19337 PR-19688

If a file system is mounted from a CDC4680 to a CONVEX machine with trailers enabled on either system, file access problems occur on files containing approximately 923 bytes. File accessing programs, such as "cat", fail with "RPC timed out" messages. Files larger than 925 bytes and smaller than 922 are handled without any problems.

PR-21250

The NFS client code can become completely stuck if one NFS server goes down. This is due to the NFS client code queuing requests for biods even when no biods are waiting for work. The NFS client code will never execute a read or write request in the application process context. When close(s) are executed, they can become stuck waiting on buffers to be flushed which are on the biod queue.

PR-21809

While running a user application, a file that is on an NFS file system is sometimes reported to be smaller than it actually is on the local (server) disk.

PR-21966

Due to the long latency times for local file access introduced by file migration systems, it is possible for NFS client and server code to hang during interaction with a file migration system.

PR-22652

The DES authentication table for nfs/rpc in the kernel can become corrupted resulting in a system panic.

---

## **refer**

PR-17602

On page 3-3 of the NFS User's Guide, the paragraph which starts: "Note the convexs and eclipse subdirectories" there is no eclipse subdirectory in the example listing. There is a "tinman" subdirectory, however.

---

## **rex**

PR-07530

If one kills a stopped rex process, the rexd on the remote host doesn't know that this has happened and the rexd will not properly clean up the executing processes.

PR-17534

Rex will hang in select once the user exits from a remote interactive rex session. The remote shell and rexd exit normally.

---

## **rpc**

PR-20297

If get\_myaddress() is unable to open a socket file descriptor, it syslog an error message and exits. It should not exit.

---

## **rstat**

PR-18576

After a power failure, there can be a problem with statd on machines all starting at the same time. The statd process will try to contact all known machines and not listen for any incoming calls. The problem is that statd will not listen to requests while they are trying to contact other statds. When many of them start up at one time, the result is many processes talking and few listening.

Workaround: This can be worked around by removing the contents of the /etc/sm and the /etc/sm.bak directories. This, however, breaks any outstanding file locks.

---

## smg

PR-20320

There is no discussion of the publickey, netid, and rpc maps in the NIS(YP) Chapter. All three are in the /etc/yp/Makefile.

PR-20321

There is no explanation in the System Manager's Guide on how the number of nfsd and biod daemons can change the performance of NFS.

PR-20580 PR-20675

The NIS chapter should mention that for sendmail to access aliases in the NIS mail.aliases maps, the sendmail O option must be set to p.

PR-21346

On page 3-19 of the Convex NFS System Manager's Guide, Using the Automounter (Order No. DSW-113), the manual should include the following error messages: "Domain name not set. Cannot use YP." "YP bind failed: can't communicate with yp-bind".

PR-22012

The following error message is not documented in the NFS System Manager's Guide. mount: serverhost:/exported/fs server not responding: RPC: Authentication error; why = Invalid client credential mount: giving up on: /exports/fs

PR-22056

The NFS Manager's Guide should mention somewhere the necessity for the nobody entry (nobody:\*:65534:65534::/:) in the /etc/passwd file.

PR-22827

A new log message has been added to the automounter which should be documented in the NFS System Manager's Guide. The message is:

path: <pathname> does not exist on host: <hostname>

Where <pathname> is the pathname that was mounted from the remote system (including any subdirectories from indirect maps using ':') and <hostname> is the name of the NFS server. To correct the problem, create the specified directory on the NFS server system.

PR-23198

The NFS System Manager's Guide should document the procedure for getting mail to work with the YP aliases map.

PR-23199

The procedure for changing the netgroup map, then getting mount to recognize the new netgroup entry should be added to the NFS smg.

---

## **YP**

PR-15777

If a system configured to use YP is booted at a time when no YP server is available, the boot does not complete and nothing can be done to the system short of a reboot.

PR-22797

We have two convex machines. One (A) is the YP master. It is running ypserv, ypbind and yppasswdd. The second machine (B) is a YP slave. it is running ypbind and yppasswdd. Reading maps on both machines works fine. However, problems arise when a user wants to change his password. on (A) password checking is done against /etc/passwd, not against the YP map. on (B) password checking is done against the YP map. Changing a password on (A) with /bin/passwd does not change the YP map.

PR-23200

Attempting to build yp maps for a non-existent domain will result in changes to the maps in the default domain.







