

Serviceguard NFS Toolkit for Linux Version A.01.03 Release Notes



i n v e n t

Manufacturing Part Number : T1442-90012

July 2005

Legal Notices

The information in this document is subject to change without notice.

Hewlett-Packard makes no warranty of any kind with regard to this manual, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Hewlett-Packard shall not be held liable for errors contained herein or direct, indirect, special, incidental or consequential damages in connection with the furnishing, performance, or use of this material.

Warranty. A copy of the specific warranty terms applicable to your Hewlett-Packard product and replacement parts can be obtained from your local Sales and Service Office.

Restricted Rights Legend. Use, duplication or disclosure by the U.S. Government is subject to restrictions as set forth in subparagraph (c) (1) (ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.227-7013 for DOD agencies, and subparagraphs (c) (1) and (c) (2) of the Commercial Computer Software Restricted Rights clause at FAR 52.227-19 for other agencies.

Hewlett-Packard Company
19420 Homestead Road
Cupertino, California 95014 U.S.A.

Use of this manual and flexible disk(s) or tape cartridge(s) supplied for this pack is restricted to this product only. Additional copies of the programs may be made for security and back-up purposes only. Resale of the programs in their present form or with alterations, is expressly prohibited.

Copyright Notices

© Copyright 2001-2005 Hewlett-Packard Development Company, L. P., all rights reserved.

Reproduction, adaptation, or translation of this document without prior written permission is prohibited, except as allowed under copyright laws.

© Copyright 1979, 1980, 1983, 1985-96 Regents of the University of California. This software is based in part on the Fourth Berkeley Software Distribution under license from the regents of the University of California.

© Copyright 1986-1997 Sun Microsystems, Inc.

Trademark Notices

Serviceguard® is a registered trademark of Hewlett-Packard Company, and is protected by copyright.

NFS® is a registered trademark of Sun Microsystems, Inc.

Red Hat® is a registered trademark of Red Hat Software, Inc.

Linux® is a registered trademark of Linus Torvalds.

UNIX® is a registered trademark of The Open Group.

1 Serviceguard NFS for Linux Version A.01.03 Release Notes

Announcements

Serviceguard NFS for Linux is a separate set of shell scripts, and a binary file that allow you to configure Serviceguard packages that mount highly available networked file systems.

An NFS server is a host that “exports” its local directories (makes them available for client hosts to mount using NFS). On the NFS client, these mounted directories look to users like part of the client’s local file system. With Serviceguard NFS, the NFS server package containing the exported file systems can move to a different node in the cluster in the event of failure. After Serviceguard starts the NFS package on the adoptive node, the NFS file systems are re-exported from the adoptive node with minimum disruption of service to users. The client side “hangs” until the NFS server package comes up on the adoptive node. When the service returns, the user can continue access to the file. You do not need to restart the client.

The following version of the NFS Toolkit is now being made available:

- Product T1442 — version A.01.03 — software and license

Serviceguard NFS for Linux is being released for use with the Linux operating system running Serviceguard A.11.14.02 and beyond. Support is provided for the following platforms:

Serviceguard A.11.14.02:

- HP Proliant Servers
- Red Hat 7.3
- RedHat AdvanceServer 2.1
- SuSE Linux Enterprise Server 8 powered by UnitedLinux 1.0 (subsequently referred to as SLES 8/UL 1.0)

Serviceguard A.11.15:

Announcements

- HP Integrity Servers and HP ProLiant Servers
- SuSE Linux Enterprise Server 8 powered by UnitedLinux 1.0
- RedHat Enterprise Server 3

Serviceguard A.11.16:

- HP Integrity Servers, HP ProLiant Servers and x86 Systems
- SuSE Linux Enterprise Server 8 powered by UnitedLinux 1.0
- RedHat Enterprise Server 3
- RedHat Enterprise Server 4
- Novell/SuSE Linux Enterprise Server 9 (subsequently referred to as SLES 9)

Complete NFS Toolkit product documentation is provided in the manual *Managing Serviceguard NFS for Linux*.

What's in this Version

The Serviceguard NFS Toolkit for Linux provides high availability for an NFS server application. The A.01.03 version of Serviceguard NFS Toolkit for Linux contains the following:

- Simple failover from an active NFS server node to an idle NFS server node.
- Failover from one active NFS server node to another active NFS server node, where the adoptive node supports more than one NFS package after the failover.
- A host configured as an adoptive node for more than one NFS package. The host may also be prevented from adopting more than one failed package at a time.
- Cascading failover, where a package may have several adoptive nodes configured to run the package.
- Fixes

One shell script (the NFS control script) is provided as a template can be customized for your specific needs.

One configuration file (`hanfs.conf`) is introduced which allows the option to monitor the quotad daemon. If `QUOTA_MON` is set to `YES`, toolkit will start `rpc.quotad` (if not already running) and will monitor this daemon. If it is set to `NO`, the toolkit will not start or monitor `rpc.quotad`. The default value for `QUOTA_MON` variable is `YES`.

After the shell scripts are installed they are located in the following location:

- RedHat distributions: `/usr/local/cmcluster/nfstoolkit`
- SLES distributions: `/opt/cmcluster/nfstoolkit`

The binary file is located in `/usr/bin` on your Linux platforms.

Complete details about supported configurations are found in the *Serviceguard for Linux Cluster Configuration Guide*, available on the web at:

<http://www.hp.com/servers/proliant/highavailability/serviceguard>

What Manuals are Available for This Version

The following manual containing information about HA NFS is included with Serviceguard NFS for Linux A.01.03:

- *Managing Serviceguard NFS for Linux* (HP Part Number T1442-90008).

Online versions of the user's guides and white papers are also available on Hewlett-Packard's Linux web page:

<http://docs.hp.com/linux/>

For information about configuring packages using the toolkits, refer to the following manuals, which are shipped with *Serviceguard for Linux* version number A.11.15 and *Serviceguard NFS for Linux*. In addition, more updated versions may be available on

<http://docs.hp.com/hpux/ha>.

Also, be sure to review the README file that accompanies the toolkit you are using.

- *Serviceguard NFS for Linux*
- *Managing Serviceguard for Linux*

Further Information

Support information, including current information on patches and known problems, is available from Hewlett-Packard IT center:

<http://itrc.hp.com> (Americas and Asia Pacific)

<http://europe.itrc.hp.com> (Europe)

The most recent versions of user's guides, release notes, and white papers are available on Hewlett-Packard's high availability documentation web page:

<http://docs.hp.com/hpux/ha>

For linux documentation:

<http://docs.hp.com/linux/>

Additional information about Serviceguard and related high availability topics may be found on Hewlett-Packard's web page:

<http://www.hp.com/go/ha>

The NFS Toolkit for Linux is not available in native language versions.

Compatibility Information and Installation Requirements

In addition to the instructions provided below, see *Managing Serviceguard NFS on Linux* (T1142-90007) for more detailed installing and configuring instructions.

IMPORTANT

You cannot use Serviceguard NFS Toolkit without NFS Services. The NFS server programs must be installed on your Linux system before you install, configure, and test your NFS package. (When installing the Linux distribution, you select the “NFS Server” option. The installation program will install both NFS kernel and utility.)

Installing Serviceguard NFS for Linux

1. Before you start, use the following steps to remove any previous version of Serviceguard NFS for Linux:

To query the NFS Toolkit from the rpm database use:

```
# rpm -qa |grep nfstoolkit
```

If any part of the NFS Toolkit is installed, erase it using:

```
# rpm -e nfstoolkit-<release>
```

2. To install the NFS Toolkit use the following commands:

RedHat:

```
# rpm -i nfstoolkit-A.01.03-0.product.redhat.i386.rpm
```

The files will be installed in the
/usr/local/cmcluster/nfstoolkit and /usr/lib directories.

SLES Distributions:

```
# rpm -I nfstoolkit-A.01.03-0.product.suse.i386.rpm
```

The files will be installed in the /opt/cmcluster/nfstoolkit and
/usr/lib directories.

The following files are part of the toolkit:

- README. Description of the tool kit contents.
- hanfs.sh. The NFS control script template that starts and stops NFS daemons and exports and unexports file systems.
- nfs.mon. The NFS monitor script.
- /usr/bin/sync_rmtab. Remote mount table synchronization binary code.
- toolkit.sh: The interface script between the Package Control Script and hanfs.sh.
- hanfs.conf: The nfs configuration file that specifies whether to monitor quotad daemon.

NOTE

In the following steps, <dir> refers to the directory /usr/local for RedHat environments, and /opt for SLES environments.

3. Run `cmmakepkg` command to generate a package configuration file and package control script template to the <dir>/cmcluster/nfstoolkit directory with the following :

```
# cd <dir>/cmcluster/nfstoolkit
# cmmakepkg -p pkg.conf
# cmmakepkg -s pkg.cnt1
```

4. Create a directory for your package files called <dir>/cmcluster/<pkg_name>
5. Issue the following command to copy the Serviceguard NFS template files to the newly created package directory:

```
# cp <dir>/cmcluster/nfstoolkit/* \
  <dir>/cmcluster/<pkg_name>
```

Copying the Template Files

If you run only one Serviceguard NFS package in your Serviceguard cluster, you do not have to copy the template files. However, if you will run multiple Serviceguard NFS packages, each package must have its own package directory, package configuration file and control scripts. For each Serviceguard NFS package run, make a copy of all the package files including package configuration file (`pkg.conf`), package control

Compatibility Information and Installation Requirements

script (`pkg.cntl`), interface script (`toolkit.sh`), NFS Control Script (`hanfs.sh`), NFS configuration file (`hanfs.conf`) and NFS monitor script (`nfs.mon`). Use package specific file names for the scripts, such as `pkg1.conf` and `pkg1.cntl`.

NOTE

`pkg.cntl`, `toolkit.sh`, `hanfs.sh` and `hanfs.conf` should be in the same directory. Do not rename the `toolkit.sh`, `hanfs.sh`, `hanfs.conf` and `nfs.mon`. These files are hard coded in the two control scripts.

Known Problems and Workarounds

The following describes known problems with the NFS Toolkit and workarounds for them. However, this is subject to change without notice. For the most current information contact your HP support representative.

More recent information on known problems and workarounds may be available on the Hewlett Packard IT Resource Center:

<http://itrc.hp.com> (Americas and Asia Pacific)

<http://europe.itrc.hp.com> (Europe)

JAGae55739: HA NFS and ‘Stale NFS Handle

What is the problem?

Serviceguard relies on LVM to manage its shared storage containing data and files of applications managed by SG. Just as with other applications, an NFS instance managed by SG means that the specific instance is “active” on one node at a time, with all resources available to that node only (including volume groups configured as a resource for the NFS package). If the package goes down on node 1, the resources are released so that the second node can “claim” the resources. The package is brought up, and the instance is now “active” on node 2. NFS clients continue to connect to the server, unaware that the server has migrated from one node to another.

Behavior has changed in Linux kernel 2.6 such that the kernel uses lvm2 with the device mapper to virtualize devices instead of using the actual physical names. In this implementation, the actual device node for a logical volume is dynamically created upon vg activation, meaning a vg that starts out on node 1 but is failed over to node 2 can very easily end up with a different minor number after fail over. This will result in clients who connected before the failover getting a “stale NFS handle”. The volume groups are created with dynamic minor numbers. This causes NFS problems (‘stale nfs handle’) on the client when the NFS package is migrated to another node

What is the workaround? There are two ways to address this issue:

- Create the logical volumes with persistent minor numbers.

- Export the filesystem with an assigned filesystem identification.

NOTE

Refer to the README File for more detailed information.

A new configuration file called `hanfs.conf` has been introduced in this version, which provides the choice of monitoring the `rpc.rquotad` daemon (which is an optional daemon).

JAGad91624: NFS remote mount table discrepancy on package fail over

What is the problem? The counter number of remote mount entries in the `/var/lib/nfs/rmtab` may or may not be correct after running NFS packages for a long period. This problem does not disrupt any client access of NFS files.

Each entry in the `/var/lib/nfs/rmtab` file contains three fields. These fields are: Client IP address, Exported file system, and Counter number of the exported file system mounted by a client.

An `rmtab` file, for example, may contain the following entries:

```
128.12.148.2:/expo/pkg1fs:0x00000001
128.12.148.3:/expo/pkg1fs:0x00000001
128.12.148.2:/expo/pkg2fs:0x00000001
```

The counter field, in normal case, contains `0x00000001` (e.g. the counter can be 2 if the same client mounted the same file system twice). The counter may or may not be updated correctly when a client un-mounts the file system.

What is the workaround?

The system administrator needs to be aware of the counter number that may not be consistent with the number of clients currently mounting the file systems. If the file, `/var/lib/nfs/rmtab`, or the counter number becomes too large, take the following steps in order to maintain the file:

1. Verify that clients are not currently mounting the file systems.
2. Verify that all NFS packages are running on all nodes in the cluster.
3. On each node, perform:

```
# cat /dev/null > /var/lib/nfs/rmtab
```

4. Resume the normal client mounting process.

Patches and Fixes in this Version

Patches

There are no known patches at the time of this publication. However, this is subject to change without notice. For the most current information contact your HP support representative.

NOTE

Patches can be superseded or withdrawn at any time. Be sure to check the status of any patch before downloading it.

An updated list of patches is available on the Hewlett Packard IT Resource Center.

<http://itrc.hp.com> (Americas and Asia Pacific)

<http://europe.itrc.hp.com> (Europe)

Fixes

The following defects are fixed in version A.01.03 of the HANFS toolkit. This is subject to change without notice. For the most current information contact your HP support representative.

Defect Number Problem and Resolution

(no JAG avail.) *Problem:* “ps issue”: The Linux command “ps pid” will sometimes return empty, even when the pid actually exists. `nfs.mon` uses this command. If this command incorrectly reports that monitored processes are not running, the application fails over from node 1 to node 2.

Resolution: `nfs.mon` has been enhanced to remove the dependency on `ps`. Instead of executing `ps`, `nfs.mon` has been modified to use `/proc` to detect processes.

NOTE

Refer to the README File for more detailed information.

Software Availability in Native Languages

The NFS Toolkit for Linux is not available in native language versions.