

# **Serviceguard NFS Toolkit for Linux Version A.01.04 Release Notes**



**Manufacturing Part Number : T1442-90013  
September 2006**

---

## Legal Notices

© Copyright 2006 Hewlett-Packard Development Company, L.P.

Publication Date: 2006

Confidential computer software. Valid license from HP required for possession, use, or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

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

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

NIS™ is a trademark of Sun Microsystems, Inc.

UNIX® is a registered trademark of The Open Group.

Linux® is a registered trademark of Linus Torvalds.

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

# 1 Serviceguard NFS for Linux Version A.01.04 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.01.04 — software and license

Serviceguard NFS for Linux is being released for use with the Linux operating system running Serviceguard A.11.16 and beyond. Support is provided for the following platforms:

### Serviceguard A.11.16:

- HP Integrity Servers, HP ProLiant Servers and x86 Systems
- RedHat Enterprise Server 3
- RedHat Enterprise Server 4
- Novell/SuSE Linux Enterprise Server 9 (subsequently referred to as SLES 9)
- SLES 10

**Announcements**

---

**NOTE**

At the time of publication, the NFS toolkit version A.01.04 supports Serviceguard version A.11.16 and later.

---

For more details on Serviceguard versions and corresponding supported Linux distro versions, check the Serviceguard's certification matrix located at <http://www.hp.com/info/sglx>

Complete NFS Toolkit product documentation is provided in the *Managing Serviceguard NFS for Linux* user's guide.

---

## What's in this Version

The Serviceguard NFS Toolkit for Linux provides high availability for an NFS server application. The A.01.04 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.

After the NFS toolkit is installed the shell scripts 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 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 contains information about HA NFS and is included with Serviceguard NFS for Linux VA.01.04:

- *Managing Serviceguard NFS for Linux*

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

## Additional Reading

The following document contains additional useful information:

- *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 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* 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:*

Proliant Servers:

```
# rpm -i nfs-toolkit-A.01.04-0.product.redhat.i386.rpm
```

Integrity Servers:

```
# rpm -i nfs-toolkit-A.01.04-0.product.redhat.ia64.rpm
```

The files will be installed in the

`/usr/local/cmcluster/nfstoolkit` and `/usr/lib` directories.

*SLES Distributions:*

Proliant Servers:

```
# rpm -i nfs-toolkit-A.01.04-0.product.suse.i386.rpm
```

Integrity Servers:

```
# rpm -i nfs-toolkit-A.01.04-0.product.suse.ia64.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 toolkit 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 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`, `nfs.mon`, 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)

### **JAGaf57739: 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).

### **JAGag06739: Serviceguard/LX NFS server returns ESTALE when package is brought down**

**What is the problem?** When the NFS package is brought down, the client process gets ESTALE from the server while it should not. There is a delay when the interface is brought down and that package could still reach the NFS server layer.

This is due to `exportfs -u`, and the ESTALE is returned. When an interface is brought down, the routes do not get flushed immediately. Therefore, the server keeps temporarily responding for after shutting down the interface.

#### **What is the workaround?**

Do the following steps on all the nodes that are configured for running NFS package to set `min_delay` to zero while the system is running:

1. Run the command  

```
# echo 0 > /proc/sys/net/ipv4/route/min_delay
```
2. Add the following line in `/etc/sysctl.conf` (in order to have persistent `min_delay` value across server boots)  

```
net.ipv4.route.min_delay = 0.
```

---

**NOTE**

---

Instructions in steps 1 and 2 needs to be done all nodes.

Instruction in step 2 ensures that `min_delay` is set to zero after system reboot. This ensures that persistent `min_delay` value is set across server reboots and hence after reboot, step 1 need not be followed.

---

## **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**

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

---

## **Software Availability in Native Languages**

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

