

Tru64 UNIX Best Practice

Providing a Cluster with Highly Available NFS-mounted File Systems via CDSLs

July 2001

**Product Version: TruCluster Server Version 5.0A and
later**

This Best Practice describes how to use context-dependent symbolic links (CDSLs) to make a Network File System (NFS) file system that is mounted on cluster members highly available in TruCluster™ Server Version 5.0A and later.

Contents

Providing a Cluster with Highly Available NFS-mounted File Systems via CDSLs

Is This Best Practice Right for You?	2
Before You Begin	2
Applying the Best Practice	3
Verifying Success	3
Troubleshooting	3
Alternative Practices	4
Comments and Questions	4
Legal Notice	5

Providing a Cluster with Highly Available NFS-mounted File Systems via CDSLs

This Best Practice describes how to use context-dependent symbolic links (CDSLs) to make a Network File System (NFS) file system that is mounted on cluster members highly available in TruCluster™ Server Version 5.0A and later.

When you mount an NFS file system on a clusterwide directory, the member where you mount the NFS file system acts as the Cluster File System (CFS) server of the file system to the other cluster members. I/O activity between the cluster and the external NFS server is routed through the mounting member by CFS. CFS maintains cache coherency across cluster members, ensuring that all members have the same view of the NFS file system at all times.

If the cluster member that mounted the NFS file system goes down, the NFS mount is removed from the namespace, and applications that were using it receive file access errors. The former mount point is forcibly unmounted automatically. The file system must be remounted on another member. Files must be reopened, users must `cd` back into directories, and similar actions must be taken to return to the status that existed before the mounting member became unavailable.

By using CDSLs, you can provide high availability to NFS-mounted file systems and avoid the temporary interruption to file access that occurs when NFS file systems are served to the cluster via CFS and the CFS server becomes unavailable. However, using CDSLs as NFS mount points means the loss of the guaranteed cache coherency provided by CFS.

Without cache coherency, data consistency is maintained by NFS. The behavior will be the same as if the members were not clustered and had separately mounted the file systems.

See the Tru64 UNIX Best Practices Web page for more information about Best Practices documentation.

Is This Best Practice Right for You?

Not all Best Practices apply to all configurations, so you must be sure that it is appropriate for your system and circumstances. To use this Best Practice, you must meet the requirements described in the following table:

Requirement	Description
Operating System	TruCluster Server Version 5.0A or later running on the same version of Tru64™ UNIX.
System Configuration	A cluster consisting of two or more members. The examples in this Best Practice are based on a two-member cluster.
Impact on Availability	You can apply this Best Practice without rebooting any cluster members.

Use this Best Practice if both of the following conditions are true:

- The cluster is an NFS client and you want high availability to some of the NFS-mounted file systems.
- The cache coherency provided by CFS is not required for the NFS file systems that you want to make highly available.

Before You Begin

You must have root login privileges to apply this Best Practice.

Before applying this Best Practice, you need to understand how a CDSL works.

A CDSL is a symbolic link that contains a variable, {memb}, that identifies a cluster member. For example:

```
/etc/rc.config --> ../cluster/members/{memb}/etc/rc.config
```

During pathname resolution, the variable is resolved into a target: {memb} is replaced with the string member_n, where *n* is the member ID of the current member.

For example, on a cluster member with member ID 2, the pathname /cluster/members/{memb}/etc/rc.config resolves to /cluster/members/member2/etc/rc.config. In this way, a CDSL provides a single file name that points to one of several files, depending on which member the file is accessed from.

Applying the Best Practice

Before you use CDSLs to make an NFS file system highly available in a cluster, be sure to read *Before You Begin*.

Perform the following steps to use a CDSL as the mount point:

1. Create the mount point. This example uses `local_mnt` as the name of the mount point.

```
# mkdir /local_mnt
```

2. Use the `mkcdsl -a` command to convert the `/local_mnt` directory into a CDSL. This copies the directory to a member-specific area on each member.

```
# mkcdsl -a /local_mnt
```

3. Log in to each cluster member where you want access to the remote file system and mount the NFS file system. Use the same mount command on each member. For example,

```
# mount -t nfs /filesystem@remote_server /local_mnt
```

4. Update `/etc/fstab` to include the new mount information so that, in the future, the mount can be performed automatically on each cluster member.

Verifying Success

After you apply this Best Practice, you can verify whether it was successful. To do so, log in to each cluster member and try to access the NFS-mounted file system.

If the Best Practice was not successful, see *Troubleshooting* for information about identifying and solving problems.

Troubleshooting

If you determine that the Best Practice was not successful, as described in *Verifying Success*, use the following table to identify and solve problems:

Problem	Possible Solutions
The file system is available on some but not all cluster members.	Log in to each member where the file system is not available and mount the NFS file system, as described in step 3 in <i>Applying the Best Practice</i> .
When you attempt to mount the NFS file system, you get a message similar to: remote-server:/filesystem on /local_mnt No such file or directory	Use the <code>ls -l</code> command to verify that the CDSL was created. For example, <code>ls -l /local_mnt</code> The result should show a link with a {memb} directory in the pathname. If the CDSL does not exist, create it with the <code>mkcdsl -a</code> command, and mount the file system, as described in steps 2 and 3 in <i>Applying the Best Practice</i> .

Alternative Practices

If you require an NFS-mounted file system to have the guaranteed cache coherency across cluster members provided by CFS, do not mount the file system on a CDSL. Instead, mount it on a clusterwide directory. Use AutoFS to automatically remount the file system in the event that the mounting member becomes unavailable.

Although AutoFS takes care of unmounting the old mount points and remounting the file system on an available member, the failover is not transparent to users and to applications that were using the NFS files at the time the member mounting the NFS file system failed.

AutoFS is not available in TruCluster Server Version 5.0A. You can instead use automount, which provides functionality similar to AutoFS, but it lacks the scalability.

For more information about AutoFS, see *TruCluster Server Cluster Administration*, *Tru64 UNIX Network Administration*, `autofs(8)`, and `autofs mount(8)`.

For more information about automount, see *Tru64 UNIX Network Administration* and `automount(8)`.

Comments and Questions

We value your comments and questions on the information in this document. Please mail your comments to us at this address:

`best_practices@zk3.dec.com`

Legal Notice

Compaq, the Compaq logo, and TruCluster Registered in U.S. Patent and Trademark Office. Tru64 is a trademark of Compaq Information Technologies Group, L.P. in the United States and other countries. UNIX is a trademark of The Open Group in the United States and other countries. All other product names mentioned herein may be trademarks of their respective companies.

Confidential computer software. Valid license from Compaq 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.

Compaq shall not be liable for technical or editorial errors or omissions contained herein. The information in this document is provided "as is" without warranty of any kind and is subject to change without notice. The warranties for Compaq products are set forth in the express limited warranty statements accompanying such products. Nothing herein should be construed as constituting an additional warranty.