

TruCluster Server

Release Notes

Part Number: AA-RHH1C-TE

August 2000

Product Version: TruCluster Server Version 5.1

Operating System and Version: Tru64 UNIX Version 5.1

This manual provides important information about TruCluster™ Server Version 5.1.

© 2000 Compaq Computer Corporation

COMPAQ and the Compaq logo Registered in U.S. Patent and Trademark Office. Alpha, AlphaServer, StorageWorks, TruCluster, and Tru64 are trademarks of Compaq Information Technologies Group, L.P.

Microsoft and Windows are trademarks of Microsoft Corporation. UNIX and The Open Group are trademarks of The Open Group. All other product names mentioned herein may be trademarks or registered 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 subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND. THE ENTIRE RISK ARISING OUT OF THE USE OF THIS INFORMATION REMAINS WITH RECIPIENT. IN NO EVENT SHALL COMPAQ BE LIABLE FOR ANY DIRECT, CONSEQUENTIAL, INCIDENTAL, SPECIAL, PUNITIVE, OR OTHER DAMAGES WHATSOEVER (INCLUDING WITHOUT LIMITATION, DAMAGES FOR LOSS OF BUSINESS PROFITS, BUSINESS INTERRUPTION OR LOSS OF BUSINESS INFORMATION), EVEN IF COMPAQ HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES AND WHETHER IN AN ACTION OF CONTRACT OR TORT, INCLUDING NEGLIGENCE.

The limited warranties for Compaq products are exclusively set forth in the documentation accompanying such products. Nothing herein should be construed as constituting a further or additional warranty.

Contents

About This Manual

1 Features

2 Restrictions

2.1	CFS Restrictions	2-1
2.2	AdvFS Restriction	2-2

3 Known Problems

3.1	Firmware Notes	3-1
3.1.1	Alpha System Reference Manual Console Firmware Required for TruCluster Server	3-1
3.2	Installation Note	3-1
3.2.1	Do Not Use the Installation Branch of the Software Option of the SysMan Menu	3-1
3.3	Cluster Creation Notes	3-2
3.3.1	The clu_create Command Fails with a "File System Full" Message If /var is full	3-2
3.3.2	The clu_create Command Does Not Add the First Member's Fully Qualified Hostname to the /etc/cfgmgr.auth File	3-2
3.4	Member Addition Notes	3-2
3.4.1	"File Exists" Messages	3-2
3.4.2	What to Do When the Kernel Build of a Newly Added Member Fails	3-3
3.5	Bootting Notes	3-3
3.5.1	Member Hangs with Lock Message During Boot	3-3
3.5.2	CNX Panic During Boot	3-3
3.5.3	Bootting a New Member Without a Cluster License Displays ATTENTION Message	3-3
3.5.4	Boot All Cluster Members Before Starting Applications That Use the Memory Channel API Library	3-4
3.5.5	Memory Channel (rm) Error Messages on Two-Member Clusters with Memory Channel 2 and Virtual Hub	3-4

3.6	Cluster Shutdown Notes	3-5
3.6.1	Do Not Shut Down a Cluster Member to Single-User Mode	3-5
3.7	Member Deletion Notes	3-5
3.7.1	Deleting a Member with an Inaccessible or Bad Member Boot Disk	3-6
3.7.2	The <code>clu_delete_member</code> Command Displays Confusing Messages	3-6
3.8	File System Notes	3-7
3.8.1	Cluster Member Shutdown Unmounts NFS File Systems That Are Not Busy	3-7
3.8.2	CFS Relocation Failures Involving Applications That Wire Memory	3-7
3.8.3	The <code>cfsmgr</code> Command Dumps Core When the <code>-p</code> or <code>-m</code> Option Is Used with the <code>-F raw</code> Option	3-8
3.9	Do Not Relocate AutoFS	3-8
3.10	LSM Note	3-9
3.10.1	The <code>voldisk list</code> Command Does Not Display All Non-LSM Disks in a Cluster	3-9
3.11	CAA Notes	3-9
3.11.1	Problems with <code>caa_relocate</code> When Multiple Interdependent Applications Are Specified	3-9
3.11.2	Action Scripts PATH Considerations	3-10
3.11.3	The <code>caa_register -u</code> Command Does Not Correctly Update a Nonapplication Resource's State	3-10
3.11.4	Changes to HOSTING_MEMBERS List Can Leave Applications Running on Disallowed Members	3-10
3.11.5	CAA Fails to Correctly Manage Applications on a Member During a Rolling Upgrade	3-11
3.11.6	CAA Events Are Malformed When Viewed from the Event Manager (EVM) Viewer	3-11
3.11.7	CAA Lets You Update a Running Application's Required Resources	3-11
3.11.8	SysMan Station Shows CAA Application Resources in UNKNOWN State as Having an Error	3-12
3.11.9	<code>caa_stat -t -v</code> Output May Be Corrupted by Long Resource Names	3-12
3.12	Cluster Alias Notes	3-12
3.12.1	The <code>aliasd</code> Daemon Sets Up Duplicate Subnets	3-12
3.12.2	<code>cluamgr</code> : Cannot Enable or Disable More Than One Interface at a Time	3-13
3.13	Rolling Upgrade Notes	3-13

3.13.1	Rolling Upgrade: Do Not Shut Down a Cluster Member To Single-User Mode	3-13
3.13.2	Deleting Layered Products During a Rolling Upgrade	3-13
3.13.3	Rolling Upgrade Disk Space Checks	3-13
3.13.4	The clu_upgrade preinstall Command Ignores User Response	3-14
3.13.5	"Cannot Rename..." or "Could Not Get Property..." Messages	3-14
3.13.6	Updating of Worldwide Language Support from CD-ROM During a Rolling Upgrade	3-14
3.13.7	Do Not Install Software After Performing the clu_upgrade setup Command	3-14
3.13.8	Run Updates to dop Database Only on Upgraded Members	3-15
3.13.9	Rolling Upgrade Does Not Update OS Version in /etc/motd	3-15
3.13.10	CAA Does Not Correctly Manage Applications During a Rolling Upgrade	3-15
3.14	Miscellaneous Administration Notes	3-15
3.14.1	Disabling Internet Services on a Per-Member Basis	3-15
3.14.2	The hwmgr -show comp Command May Report an Inconsistency Error When Creating a Clusterwide Name for a SCSI Device	3-16
3.14.3	The hwmgr -scan scsi Command Does Not Work Clusterwide	3-16
3.14.4	Adding a Disk to a Running Cluster	3-17
3.14.5	A Media Changer Device Can Be Managed Only from Its DRD Server	3-17
3.14.6	The sysman -cli Command Requires the -focus Option to Remove Member-Specific NFS Configuration Values	3-17
3.14.7	Running Process Accounting on Large Clusters Can Exhaust Member Process Quotas	3-18
3.14.8	Login Refused on Member in Multiuser Mode After Performing shutdown now Command on Lead Cluster Member	3-18
3.14.9	Clearing Persistent Reservations on Disks Behind HSZ80 and HSG80 Controllers	3-18
3.15	SysMan Menu Note	3-19
3.15.1	Nonroot Users Cannot Run System Management Applications	3-19
3.16	SysMan Station Notes	3-20
3.16.1	Unable to Expand Host Object in a Cluster	3-20

3.16.2	SysMan Station Might Display Cluster Status Incorrectly	3-20
3.16.3	SysMan Station Might Display New Hardware Objects Incorrectly	3-22
3.16.4	Properties Might Not Be Displayed for Selected Objects ..	3-22
3.16.5	Some SysMan Station Applications Display Wrong Target Member Name	3-22
3.17	Documentation Notes	3-23
3.17.1	Software Installation: Incorrect No WLS and Disk Space Example	3-23

4 Usage

4.1	Managing Third-party Storage	4-1
-----	------------------------------------	-----

Index

Tables

1-1	TruCluster Server Version 5.1 Features	1-1
-----	--	-----

About This Manual

This manual provides important information about TruCluster™ Server Version 5.1, which runs on the Tru64™ UNIX Version 5.1 operating system.

Audience

TruCluster Server users, administrators, and programmers should read this manual.

Organization

This manual contains the following chapters:

- *Chapter 1* lists the special features of TruCluster Server Version 5.1.
- *Chapter 2* lists restrictions on the use of certain components in TruCluster Server Version 5.1.
- *Chapter 3* discusses known problems with TruCluster Server Version 5.1.
- *Chapter 4* provides information that was not included in the TruCluster Server manuals.

Related Documents

Consult the following manuals for assistance in cluster hardware and software configuration, installation, administration, and programming tasks:

- TruCluster Server *Software Product Description* (SPD) — The comprehensive description of the TruCluster Server Version 5.1 product. You can find the latest version of the SPD and other TruCluster Server documentation at the following URL:
http://tru64unix.compaq.com/faqs/publications/pub_page/cluster_list.html
- TruCluster Server *Technical Overview* — Provides an overview of the TruCluster Server technology.
- TruCluster Server *Hardware Configuration* — Describes how to set up the processors that are to become cluster members, and how to configure cluster shared storage.
- TruCluster Server *Software Installation* — Describes how to install the TruCluster Server software.

- TruCluster Server *Highly Available Applications* — Describes how to deploy existing applications on a TruCluster Server cluster, and how to write cluster-aware applications.
- TruCluster Server *Cluster Administration* — Describes cluster-specific administration tasks.

Note

Read the Tru64 UNIX *Release Notes* and the Tru64 UNIX *System Administration* manual to become familiar with restrictions and new features in the base operating system before you install, configure, and use your TruCluster Server cluster.

Reader's Comments

Compaq welcomes any comments and suggestions you have on this and other Tru64 UNIX manuals.

You can send your comments in the following ways:

- Fax: 603-884-0120 Attn: UBPG Publications, ZKO3-3/Y32
- Internet electronic mail: readers_comment@zk3.dec.com

A Reader's Comment form is located on your system in the following location:

```
/usr/doc/readers_comment.txt
```

- Mail:

Compaq Computer Corporation
UBPG Publications Manager
ZKO3-3/Y32
110 Spit Brook Road
Nashua, NH 03062-2698

A Reader's Comment form is located in the back of each printed manual. The form is postage paid if you mail it in the United States.

Please include the following information along with your comments:

- The full title of the book and the order number. (The order number is printed on the title page of this book and on its back cover.)
- The section numbers and page numbers of the information on which you are commenting.
- The version of Tru64 UNIX that you are using.
- If known, the type of processor that is running the Tru64 UNIX software.

The Tru64 UNIX Publications group cannot respond to system problems or technical support inquiries. Please address technical questions to your local system vendor or to the appropriate Compaq technical support office. Information provided with the software media explains how to send problem reports to Compaq.

Conventions

The following typographical conventions are used in this manual:

%

\$

A percent sign represents the C shell system prompt.
A dollar sign represents the system prompt for the Bourne, Korn, and POSIX shells.

#

A number sign represents the superuser prompt.

% **cat**

Boldface type in interactive examples indicates typed user input.

file

Italic (slanted) type indicates variable values, placeholders, and function argument names.

[|]

{ | }

In syntax definitions, brackets indicate items that are optional and braces indicate items that are required. Vertical bars separating items inside brackets or braces indicate that you choose one item from among those listed.

...

In syntax definitions, a horizontal ellipsis indicates that the preceding item can be repeated one or more times.

cat(1)

A cross-reference to a reference page includes the appropriate section number in parentheses. For example, `cat(1)` indicates that you can find information on the `cat` command in Section 1 of the reference pages.

Return

In an example, a key name enclosed in a box indicates that you press that key.

Ctrl/x

This symbol indicates that you hold down the first named key while pressing the key or mouse button that follows the slash. In examples, this key combination is enclosed in a box (for example, Ctrl/C).

1

Features

TruCluster Server Version 5.1 supports all of the features of TruCluster Server Version 5.0A plus additional enhancements that are listed in Table 1–1.

Table 1–1: TruCluster Server Version 5.1 Features

Feature	Description	For more information, see:
Support of AlphaServer™ GS320, GS160, and GS80 systems	TruCluster Server Version 5.1 supports the AlphaServer GS320, AlphaServer GS160, and AlphaServer GS80 systems.	<i>Hardware Configuration</i>
Automount scalability – support for AutoFS	AutoFS can be used as an alternative to automount and provides greater scalability in support of large clusters.	<i>Cluster Administration</i>
Rolling upgrade/rolling patch	TruCluster Server Version 5.1 supports rolling upgrades and patches. Customers who installed TruCluster Server Version 5.0A will be able to perform a rolling upgrade to TruCluster Server Version 5.1, and roll patches onto a Version 5.0A or Version 5.1 cluster.	<i>Technical Overview, Software Installation</i>
Cluster file system (CFS) concurrent direct I/O	For files opened with the <code>O_DIRECT</code> flag, TruCluster Server Version 5.1 allows remote CFS clients, as well as applications that are local to the CFS server, to read and write directly to the files. Regardless of which member originates the I/O request, direct I/O to a file does not go through the cluster interconnect to the CFS server.	<i>Cluster Administration</i>

Restrictions

This chapter lists restrictions on the use of TruCluster Server Version 5.1 features.

2.1 CFS Restrictions

The cluster file system (CFS) has the following restrictions in TruCluster Server Version 5.1:

- CFS supports the Network File System (NFS) client for read/write access.

When a file system is NFS-mounted in a cluster, CFS makes it available for read/write access from all cluster members. The member that has actually mounted it serves the file system to other cluster members.

If the member that has mounted the NFS file system shuts down or fails, the file system is automatically unmounted and CFS begins to clean up the mount points. During the cleanup process, members that access these mount points may see various types of behavior, depending upon how far the cleanup has progressed:

- If members still have files open on that file system, their writes will be sent to a local cache instead of to the actual NFS-mounted file system.
- After all of the files on that file system have been closed, attempts to open a file on that file system will fail with an `EIO` error until the file system is remounted. Applications may encounter "Stale NFS handle" messages. This is normal behavior on a standalone system, as well as in a cluster.

Until the CFS cleanup is complete, members may still be able to create new files at the NFS file system's local mount point (or in any directories that were created locally beneath that mount point).

An NFS file system does not automatically fail over to another cluster member. Rather, you must manually remount it — on the same mount point or another — from another cluster member to make it available again. Alternatively, booting a cluster member will remount those file systems that are listed in the `/etc/fstab` file that are not currently mounted and served in the cluster. (If you are using AutoFS or automount, the remount will happen automatically.)

- CFS does not support the Memory File System (MFS).

- CFS does not support a clusterwide `/proc` file system. A `/proc` file system can be mounted and accessed only on the local member.
- CFS does not support a clusterwide File-on-File Mount (FFM) file system. An FFM file system can be mounted and accessed only on the local member.
- CFS does not support clusterwide named pipes. The reader and writer of a named pipe must reside on the same member.

2.2 AdvFS Restriction

The following restriction applies to AdvFS in a cluster:

- You cannot use the `addvol` command to add volumes to a member's root domain (the a partition on the member's boot disk). Instead, you must delete the member from the cluster, use `diskconfig` or `SysMan` to configure the disk appropriately, and then add the member back into the cluster. For the configuration requirements for a member boot disk, see the TruCluster Server *Software Installation* manual.

Known Problems

The following sections describe known problems with TruCluster Server Version 5.1.

3.1 Firmware Notes

The following notes apply to firmware requirements.

3.1.1 Alpha System Reference Manual Console Firmware Required for TruCluster Server

Systems that will run TruCluster Server Version 5.1 or later require Alpha™ System Reference Manual (SRM) console firmware Version 5.7 or later.

The following table lists the minimum required firmware for specific configurations:

Minimum Firmware Required	Configuration
Version 5.7	When using HSZ80, HSG80, or HSG60 devices as member boot disks. If the cluster member is using Version 5.6 firmware or earlier, the member may fail to boot, indicating "Reservation Conflict" errors.
Version 5.6	When using <i>both</i> unit partitioning and multipathing on an HSG80 device.

You can obtain Alpha SRM console firmware from the latest Alpha Systems Firmware Update CD-ROM or the following Web location:

<http://ftp.digital.com/pub/DEC/Alpha/firmware>

3.2 Installation Note

The following note applies to installation.

3.2.1 Do Not Use the Installation Branch of the Software Option of the SysMan Menu

The Installation branch of the Software menu of the SysMan Menu application is not supported in a cluster. Use the mechanisms for installing

and deinstalling the TruCluster Server product and layered product software that are discussed in the TruCluster Server *Software Installation* and *Cluster Administration* manuals.

3.3 Cluster Creation Notes

The following notes apply to cluster creation.

3.3.1 The `clu_create` Command Fails with a "File System Full" Message If `/var` is full

If there are too many member-specific files in `/var`, the `clu_create` command fails with a "file system full" message. If this message appears, remove unneeded member-specific files from `/var`, such as files in `/var/adm/crash` and log files in `/var/adm`.

3.3.2 The `clu_create` Command Does Not Add the First Member's Fully Qualified Hostname to the `/etc/cfgmgr.auth` File

The `clu_create` command fails to add the first member's fully qualified hostname to the `/etc/cfgmgr.auth` file. The `clu_add_member` command, however, does add subsequent members' hostnames to the file.

To avoid problems with remote kernel configuration management in a cluster, manually add the first member's fully qualified hostname to the `/etc/cfgmgr.auth` file. For example:

```
member1.zk3.dec.com
member2.zk3.dec.com
member3.zk3.dec.com
```

3.4 Member Addition Notes

The following notes apply to adding new members to the cluster.

3.4.1 "File Exists" Messages

When using the `clu_add_member` command to add a cluster member, you may see the following messages. You can ignore these messages.

```
Creating cluster member-specific files:
Creating new member's root member-specific files
Creating new member's usr member-specific files
Creating new member's var member-specific files
/var/cluster/members/member5/im/webagent/IMAGES/images: File exists
/var/cluster/members/member5/im/webagent/cqmghost/images/images: File exists
/var/cluster/members/member5/im/webagent/cqmgncs/images/images: File exists
/var/cluster/members/member5/im/webagent/cqmgstor/images/images: File exists
/var/cluster/members/member5/im/webagent/cqmgstor/images/images: File exists
Creating new member's boot member-specific files
```

3.4.2 What to Do When the Kernel Build of a Newly Added Member Fails

When you first boot a newly added member, the new kernel for the member is built. If the kernel build fails, you are asked if you want to build a kernel without the optional kernel options.

Answer *no* to this question. A kernel that is built without the optional kernel options is incompatible with the existing cluster and will not boot.

If the kernel build fails for a newly added member, take steps to determine why the build failed and, after you diagnose and fix the problem, use the `doconfig` command to rebuild the kernel.

3.5 Booting Notes

The following notes discuss requirements and restrictions for booting members into a cluster.

3.5.1 Member Hangs with Lock Message During Boot

When multiple members are booted after a quorum loss situation, one member might hang and display the following message:

```
WARNING: Retrying to lock the boot partition device
```

To fix the problem, halt the hung member and reboot. The other nodes are not affected.

3.5.2 CNX Panic During Boot

During boot, a member may panic and display the following message:

```
CNX MGR: Invalid configuration for cluster seq disk
```

If this occurs, reboot the member.

3.5.3 Booting a New Member Without a Cluster License Displays ATTENTION Message

When you boot a newly added member, the `clu_check_config` utility performs a series of configuration checks. If you have not yet installed the TruCluster Server license, the TCS-UA product authorization key (PAK), on the member, the boot procedure will display:

```
Starting Cluster Configuration Check...
The boottime cluster check found a potential problem.
For details search for !!!!!ATTENTION!!!! in /cluster/admin/clu_check_log_hostname
check_cdsl_config : Boot Mode : Running /usr/sbin/cdslinvchk in the background
check_cdsl_config : Results can be found in : /var/adm/cdsl_check_list
clu_check_config : no configuration errors or warnings were detected
```

When you inspect the `/cluster/admin/clu_check_log_hostname` file, the following message is displayed:

```
/usr/sbin/caad is NOT_RUNNING !!!!!ATTENTION!!!!!
```

When the TruCluster Server license is not configured on a member, the cluster application availability daemon (`caad`) is not automatically started on that member. This is normal and expected behavior.

If you did not configure the license from within `clu_add_member` when you added the new member (as discussed in the TruCluster Server *Software Installation* manual), you can configure it later using the `lmf register` command. After the license has been installed, you can start the CAA daemon on that member using the `usr/sbin/caad` command.

3.5.4 Boot All Cluster Members Before Starting Applications That Use the Memory Channel API Library

If there is insufficient Memory Channel address space in a cluster, a booting node may have problems joining the cluster. When this is the case, one or more members may panic with an assertion failure (ICS MCT Assertion Failed), or the booting member may hang early in its boot.

Memory Channel resources are dynamically allocated as new members join the cluster. Running applications that call the Memory Channel application programming interface (API) library functions can consume required Memory Channel resources, and prevent a member from getting the resources it needs to join the cluster. To avoid this problem, boot all cluster members before starting any applications that call the Memory Channel API library functions.

3.5.5 Memory Channel (rm) Error Messages on Two-Member Clusters with Memory Channel 2 and Virtual Hub

On a two-member cluster that uses the CCMAB (MC2) Memory Channel PCI adapter configured in virtual hub mode, `rm` error messages like the following are generated when a member boots or shuts down:

```
rm error: mchan0 error_type = 0x60000000 error_count = 0xc  
time = 0x74c2c3992c4eb  
mcerr = 0x12020008 lcsr = 0xc0ff mcport = 0xc6000000
```

These messages are normal for this configuration, and you can safely disregard them.

3.6 Cluster Shutdown Notes

The following section discusses the recommended procedure for shutting down a cluster member.

3.6.1 Do Not Shut Down a Cluster Member to Single-User Mode

To take a cluster member to single-user mode, use the `shutdown -h` to halt the member, and then boot the member to single-user mode. When the system reaches single-user mode, run the `init s`, `bcheckrc`, `kloadsrv`, and `lmf reset` commands. For example:

```
# /sbin/shutdown -h now

>>> boot -fl s

# /sbin/init s
# /sbin/bcheckrc
# /sbin/kloadsrv
# /usr/sbin/lmf reset
```

Before halting a cluster member, make sure that the cluster can maintain quorum without the member's vote.

We recommend this halt and boot method because it ensures that the member will provide a minimal set of services to the cluster, and conversely, that the running cluster will have minimal reliance on the member in single-user mode.

Note

This halt and boot method contradicts the information in the TruCluster Server *Software Installation* manual, which clearly states in the rolling upgrade procedure that a member **must** be shut down to single-user mode from multiuser mode.

The reason for the change is that shutting a member down to single-user mode can take a long time, and in some cases, services such as `automount`, `autofs`, or NIS might appear to be working from the point of view of other cluster members, when in fact the services are either down or partially up.

3.7 Member Deletion Notes

The following sections discuss problems with the `clu_delete_member` command.

3.7.1 Deleting a Member with an Inaccessible or Bad Member Boot Disk

The `clu_delete_member -f` command can delete a member with an inaccessible or bad boot disk. However, be aware of the following issues:

- When you delete a member with an inaccessible boot disk, `clu_delete_member -f` does not adjust expected votes in the running cluster (as a normal `clu_delete_member` command does). If the deleted member was a voting member, use the `clu_quorum -e` command after the member has been deleted to adjust expected votes appropriately.
- The `clu_delete_member -f` command cannot delete a member if the disk label on its boot disk has been overwritten or corrupted.

To work around this problem, you must manually edit the `label:` field on the boot disk so that it contains a string of the form `clu_membern`, where *n* is the cluster member's member ID. You can then enter the `clu_delete_member` command to delete the member from the cluster.

For example:

```
% disklabel -e dsk6
# /dev/rdisk/dsk6c:
type: SCSI
disk: RZ28
label: clu_member5
flags:
bytes/sector: 512
sectors/track: 99
tracks/cylinder: 16
sectors/cylinder: 1584
cylinders: 2595
sectors/unit: 4110480
rpm: 5400
interleave: 1
trackskew: 6
cylinderskew: 30
headswitch: 0 # milliseconds
track-to-track seek: 0 # milliseconds
drivedata: 0

8 partitions:
# size offset fstype [fsize bsize cpg] # NOTE:
values n
ot exact
write new label? [y]:
re-edit the label? [n]:
```

3.7.2 The `clu_delete_member` Command Displays Confusing Messages

The `clu_delete_member` command requires that the member you intend to delete be halted at the time you enter the command. If you specify an active member to the `clu_delete_member -m memberid` command, the command should display the following message:

```
Error: Cluster member ID memberid is UP.  
Members with a state of UP cannot be deleted
```

Instead, it displays the following message:

```
Error: clu_delete_member: Bad configuration
```

To continue, halt the member that you intend to delete.

Additionally, as `clu_delete_member` proceeds to delete a member, it displays NTP peer messages and a few extraneous mail messages, such as the following:

```
Initial cluster deletion successful, member '3' can no longer join the  
cluster. Deletion continuing with cleanup.  
expected_votes: reconfigured  
Removing cluster interconnect interface 'pepicelli-mc0' from /.rhosts  
Configuring Network Time Protocol for new member  
Deleting interface 'pepperoni.zk3.dec.com' as an NTP peer  
    to member 'pepperoni.zk3.dec.com'  
Deleting interface 'polishham.zk3.dec.com' as an NTP peer  
    to member 'polishham.zk3.dec.com'  
Clusterizing mail...  
Mail has not been configured for the cluster.  
Please run /usr/sbin/mailconfig or /usr/sbin/mailsetup  
to configure mail in the cluster.
```

You can safely ignore these messages.

3.8 File System Notes

The following sections discuss issues with the CFS, AdvFS, and NFS file systems in a cluster.

3.8.1 Cluster Member Shutdown Unmounts NFS File Systems That Are Not Busy

When a cluster member shuts down (using the `shutdown` command), the `/sbin/rc0` run-level transition script unmounts all NFS file systems in the cluster that are not in active use. To provide automatic remounting of the file systems, use `AutoFS` or `automount`.

3.8.2 CFS Relocation Failures Involving Applications That Wire Memory

Applications that use the `plock()` or `mlock()` system call to lock pages of physical memory can cause the `cfsmgr` command to fail when performing a manual relocation.

If the application uses `plock()`, the domain or file system that contains the application executable cannot relocate. In the case of `mlock()`, if the locked pages are associated with files, then the file systems where those files reside cannot relocate.

In the event of failure, the `cfsmgr` command returns the following message:

```
Server Relocation Failed
Failure Reason: Invalid Relocation
```

To allow the relocation to complete for the domain or file system on which the executables reside, kill the processes (if you are able to identify them) that are running the executables using the `plock()` and `mlock()` system calls.

3.8.3 The `cfsmgr` Command Dumps Core When the `-p` or `-m` Option Is Used with the `-F raw` Option

Explicitly specifying the `-p` or `-m` option on the command line with the `-F raw` option causes the `cfsmgr` command to dump core. There is no workaround.

3.9 Do Not Relocate AutoFS

Do **not** manually relocate or stop the AutoFS, regardless of the method you use (`caa_relocate`, `caa_stop/caa_start`, `/sbin/init.d/nfsmount start/stop`, or shutting down the member running AutoFS to single user mode). Doing so may result in the following problems:

- Both indirect and direct map mountpoints that were mounted on the original member are not unmounted.
- No additional mounts can occur on indirect map intercept points (for example, if `/home/jsmith` was mounted by the original AutoFS daemon, no other entries under the `/home` directory can be mounted or accessed).

If AutoFS does not properly relocate to another cluster member, perform the following steps to recover:

1. If the member originally running the AutoFS daemon is at single user mode, bring it up to multiuser mode using the `init -3` command.
2. Stop the AutoFS service by running the following commands on this member:

```
# caa_stop -f autofs
# caa_stat autofs
```

The CAA status will indicate that the AutoFS state is OFFLINE.

3. Reboot this member.
4. Start AutoFS on another cluster member that is in multiuser mode:

```
# caa_start autofs -c new_member
```

Note

This reconstructs the `automount` mountpoints based on the parameters that were previously set in the `/etc/rc.config.common` file. Therefore, it will not account for any manual invocations of `autofs` that may have taken place with arguments not already included in that file.

3.10 LSM Note

The following note discusses a known problem with using the Logical Storage Manager (LSM) in a TruCluster Server cluster.

3.10.1 The `voldisk list` Command Does Not Display All Non-LSM Disks in a Cluster

The `voldisk list` command is supposed to show disks that are configured both under LSM and not under LSM, which gives disks that are not configured under LSM a status of `unknown`. In a cluster, however, the `voldisk list` command displays only those non-LSM disks that are directly connected to the member on which the command was executed.

For disks that are configured with LSM, use the `voldisk list` command to display consistent data from all members.

3.11 CAA Notes

The following notes discuss known problems in the cluster application availability (CAA) subsystem.

3.11.1 Problems with `caa_relocate` When Multiple Interdependent Applications Are Specified

A forced relocate that is not directed to a specific cluster member and has multiple interdependent applications specified will relocate all of these applications multiple times. For example, if `app1` depends on `app2` and `app3`, then `caa_relocate -f app1 app2 app3` causes all of these applications to relocate three times.

A forced relocate directed to a specific cluster member with interdependent applications specified relocates the applications to the specified cluster member correctly, but responds with an error. For example, `caa_relocate -c member2 -f app1 app2 app3` successfully relocates all applications to `member2`, but error messages are displayed and the return code is non-zero.

If you are relocating interdependent applications, specify only *one* application, for example, `caa_relocate -f app1` or `caa_relocate -c member2 -f app1`. All dependent applications relocate with the application specified.

Also, `caa_relocate -s` fails when used with resources with dependencies. To avoid this, do the following:

1. Identify applications with dependencies by using the `caa_stat -p` command. Those applications with entries for `REQUIRED_RESOURCES` have dependencies.
2. Use the `caa_relocate -f` command to relocate the applications with dependencies.
3. Use `caa_relocate -s member1 -c member2` to relocate the other applications.

3.11.2 Action Scripts PATH Considerations

Newly generated CAA scripts do not set the `PATH` environment variable. When they are executed, the `PATH` is set to a default value `/sbin:/usr/sbin:/usr/bin`. Therefore, you must explicitly specify most path names that are used in scripts, or you must modify the resulting scripts to explicitly set the `PATH`. Action scripts generated in previous releases may have a `PATH` that includes the current directory (`.`). Because this situation may be a potential security issue, modify these scripts to remove the current directory from the path.

3.11.3 The `caa_register -u` Command Does Not Correctly Update a Nonapplication Resource's State

When you change the `SUBNET` value in a network resource's profile or the `DEVICE_NAME` value in a tape or changer resource's profile and run the `caa_register -u` command to update the CAA registry, CAA does not update the `STATE` value for the resource. To correctly update a nonapplication resource, follow these steps:

1. Unregister the resource using the `caa_unregister` command.
2. Change the `SUBNET` or `DEVICE_NAME` attribute in the resource profile.
3. Reregister the resource using the `caa_register` command.

3.11.4 Changes to `HOSTING_MEMBERS` List Can Leave Applications Running on Disallowed Members

A change to the `HOSTING_MEMBERS` list only affects future relocations and starts. If you update the `HOSTING_MEMBERS` list in the profile of an

ONLINE application resource with a restricted placement policy, make sure that the application is running on one of the cluster members in that list. If the application is not running on one of the allowed members, run the `caa_relocate` on the application after running the `caa_register -u` command.

3.11.5 CAA Fails to Correctly Manage Applications on a Member During a Rolling Upgrade

In Tru64 UNIX Version 5.0A, when a member is brought down to single-user mode, any applications that are running on the member are stopped. However, the state of these applications is incorrectly displayed as ONLINE. This problem is fixed in Version 5.1.

Part of the process of rolling a cluster member from Tru64 UNIX Version 5.0A to Version 5.1 requires that the member, while still running Version 5.0A, be brought down to single-user mode. When this happens, the CAA state error occurs. See Section 3.13.10 for a description of how to deal with this situation.

3.11.6 CAA Events Are Malformed When Viewed from the Event Manager (EVM) Viewer

The Event Manager (EVM) viewer may display malformed CAA event messages, or messages with missing information. For example, the message:

```
CAA named is transitioning from state ONLINE to state OFFLINE on skiing
```

is displayed as:

```
CAA named is transitioning from state to state skiing
```

To work around this problem, examine the messages in the `daemon.log` file for more complete information. The messages in the `daemon` log file are in a slightly different format from those that the EVM viewer displays.

3.11.7 CAA Lets You Update a Running Application's Required Resources

The `caa_register -u` command and the CAA GUI allow you to successfully update the `REQUIRED_RESOURCES` field in the profile of an ONLINE resource with the name of a resource that is OFFLINE. (It should not allow you to update this field unless the required resource is ONLINE or the resource whose profile you are modifying is OFFLINE.)

If you accidentally update the `REQUIRED_RESOURCES` field incorrectly, you must manually start the required resource or stop the updated resource to correct their states.

3.11.8 SysMan Station Shows CAA Application Resources in UNKNOWN State as Having an Error

The SysMan Station shows any application resources that are in the UNKNOWN state as having an error and does not show what member the application is UNKNOWN on. For example, an application resource named `foo` in the UNKNOWN state is shown underneath the cluster icon as `foo (error)`.

3.11.9 `caa_stat -t -v` Output May Be Corrupted by Long Resource Names

The `caa_stat -t -v` command may display corrupted information for resources that have very long names. To work around this problem, use the `caa_stat -v` command to display the output for these resources in a nontabular form.

3.12 Cluster Alias Notes

The following notes pertain to the cluster alias subsystem.

3.12.1 The `aliasd` Daemon Sets Up Duplicate Subnets

The `aliasd` daemon does not keep track of the stanzas that it issues in the `gated` configuration file for virtual subnet aliases. If more than one alias resides in the same virtual subnet, the `aliasd` daemon incorrectly issues extra stanzas for the given subnet. This will cause `gated` to exit with an error message in `/usr/adm/syslog.dated/current/daemon.log`:

```
duplicate static route
```

To work around this problem, modify the `cluamgr` commands in `/etc/clu.alias.config` to set the `virtual` flag only once. If there are many aliases in the same virtual subnet, use the `virtual=t` argument for only one of the aliases in the subnet. For example:

```
/usr/sbin/cluamgr -a selw=3,selp=1,join,alias=DEFAULTALIAS  
/usr/sbin/cluamgr -a alias=crunch,rpri=1,selw=1,selp=1,join,virtual=t  
/usr/sbin/cluamgr -a alias=research,rpri=1,selw=1,selp=1,join
```

The virtual subnet will still exist, but by eliminating the `virtual=t` argument for the second alias (alias `research` in the example), `aliasd` will not attempt to set up a duplicate subnet for it.

3.12.2 cluamgr: Cannot Enable or Disable More Than One Interface at a Time

In `cluamgr(8)`, the syntax for the `cluamgr -r enable=` and `cluamgr -r disable=` commands indicates that you can enable or disable multiple interfaces at one time.

However, you can enable or disable only one interface at a time. To enable or disable two interfaces, you must issue two separate `cluamgr` commands. For example:

```
# cluamgr -r disable=tu0
# cluamgr -r disable=tu1
```

3.13 Rolling Upgrade Notes

The following notes discuss issues with rolling upgrade.

3.13.1 Rolling Upgrade: Do Not Shut Down a Cluster Member To Single-User Mode

To take a cluster member to single-user mode, use the `shutdown -h` command to halt the member, and then boot the member to single-user mode. When the system reaches single-user mode, run the `init s`, `bcheckrc`, `kloadsrv`, and `lmf reset` commands.

For an example and more information, see the note in Section 3.6.1.

3.13.2 Deleting Layered Products During a Rolling Upgrade

If you delete a layered product subset while a cluster is in a mixed version state during a rolling upgrade, you must manually remove the `.Old.` lock and status files in the `/usr/.smdb.` directory after the subset deletion. For example, if you remove the `ASUADM500` subset from a mixed version cluster, enter the following commands:

```
# rm -f /usr/.smdb./Old..ASUADM500.lk
# rm -f /usr/.smdb./Old..ASUADM500.sts
```

3.13.3 Rolling Upgrade Disk Space Checks

The `clu_upgrade` command performs a disk space check at the start of a rolling upgrade. However, nothing prevents a cluster member from consuming disk space during a rolling upgrade, and thus creating a situation where a later stage might not have enough disk space. For example, after a cluster passes the initial check, but a large amount of disk space in `/usr` is consumed, the `installupdate` command can fail because there is not enough disk space available when it performs its checks.

Remember that disk space is dynamic. If you know that a member will be consuming disk space during a rolling upgrade, we suggest that you add additional space before you start the upgrade.

3.13.4 The `clu_upgrade preinstall` Command Ignores User Response

When running the `clu_upgrade preinstall` command, if an error is detected while checking tagged files, `clu_upgrade` will ask you if you want to continue. The `clu_upgrade` command will ignore your answer and continue. Enter `Ctrl/C` to cancel out of the preinstall stage, address the tagged file errors, and rerun the preinstall stage.

3.13.5 "Cannot Rename..." or "Could Not Get Property..." Messages

While performing an undo operation during the install stage, you may see the following messages. You can ignore these messages.

```
Do you want to continue to undo this stage of the upgrade? [yes]:
Restoring tagged files.
Cannot rename /usr/.smbd/.RollTemp...Old...Old...winstall to /
usr/.smbd/.Old...Old...winstall
Could not get property for file /usr/.smbd/.RollTemp...Old...Old...winstall
Cannot rename /usr/.smbd/.RollTemp...Old...winstall to
/usr/.smbd/.Old...winstall
Could not get property for file /usr/.smbd/.RollTemp...Old...winstall
Cannot rename /usr/.smbd/.RollTemp...winstall to
/usr/.smbd/.winstall
/usr/.smbd/.RollTemp...winstall is not a tagged file
```

3.13.6 Updating of Worldwide Language Support from CD-ROM During a Rolling Upgrade

If you remove the Worldwide Language Support (WLS) product from a cluster while a rolling upgrade is in progress, the cluster will be put into a state that will prevent the rolling upgrade from completing successfully.

Remove the WLS product prior to starting the rolling upgrade. Use the `/usr/sbin/clu_upgrade` command to determine if a rolling upgrade is in progress. See `clu_upgrade(8)` and the TruCluster Server *Software Installation* manual for additional details on how to determine if a rolling upgrade is in progress.

3.13.7 Do Not Install Software After Performing the `clu_upgrade setup` Command

The `clu_upgrade setup` command prepares a cluster for a rolling upgrade of the operating system. Do not use the `setld` command to load software onto the cluster between performing the `clu_upgrade setup` command and rolling the first cluster member to the new version. Wait until after rolling at least one member.

If you install software between performing the `clu_upgrade setup` command and rolling a cluster member to the new version, the new files will not be processed by `clu_upgrade`. As a result, when you roll the first cluster member, these new files will be overwritten. This can cause the cluster to fail in unpredictable ways.

3.13.8 Run Updates to dop Database Only on Upgraded Members

When performing a cluster rolling-upgrade, do not issue any `dop` commands (for example `dop -a`, `dop -d`, or `dop -W`) on cluster members running the older version of the operating system. Any `dop` additions or deletions made on the members running the older operating system might be lost after `dop` commands are issued from the upgraded members.

3.13.9 Rolling Upgrade Does Not Update OS Version in /etc/motd

The rolling upgrade process does not update the operating system version in `/etc/motd`.

3.13.10 CAA Does Not Correctly Manage Applications During a Rolling Upgrade

When a member running Tru64 UNIX Version 5.0A is brought down to single-user mode during a cluster rolling upgrade from Version 5.0A to Version 5.1, any applications that are running on the member are stopped, and the state of these applications is incorrectly displayed as `ONLINE`.

Before starting a rolling upgrade on a member, relocate all applications on that member to another member. Do not use the `caa_relocate -s` command to relocate the applications. Instead, relocate each application separately.

3.14 Miscellaneous Administration Notes

The following notes discuss issues with various administration tools that are used in a cluster.

3.14.1 Disabling Internet Services on a Per-Member Basis

The *Cluster Administration* manual describes a way to disable the Internet services daemon (`inetd`) on a particular member by using the `disable` keyword in the `/etc/inetd.conf.local` file. For example:

```
finger stream tcp      nowait root    disable      fingerd
```

If you add this entry, the following error appears in `/var/adm/sys-log.dated/current/daemon.log`:

```
inetd: disable: file does not exist
```

The Internet services daemon continues to run normally, and the service is not disabled. The workaround is to remove the entry for the service from the global `/etc/inetd.conf` file and add it to the member-specific `/etc/inetd.conf.local` file on the members that you want to offer the service. For more information on configuring `inetd`, see `inetd.conf(4)`.

3.14.2 The `hwmgr -show comp` Command May Report an Inconsistency Error When Creating a Clusterwide Name for a SCSI Device

After you have used the `hwmgr -edit scsi` command to create a clusterwide unique name for a SCSI device, a subsequent `hwmgr -show comp` command may report an inconsistency on the SCSI device. The inconsistency will be visible when the `hwmgr -edit scsi` command is invoked on the second and subsequent members for the same device. The inconsistency error can be ignored in this situation.

The inconsistency error can be seen by using the following command:

```
root> hwmgr -show comp -id 373 -full

HWID:  HOSTNAME  FLAGS  SERVICE  COMPONENT NAME
-----
373:   rovel-qa1  rcd-i  iomap    SCSI-WWID:ff10000b:"media_chngr"

DSF GROUP
INSTANCE GRPFLAGS  GROUPID  SUBSYSTEM  BASENAME  L1          L2
-----
0         40         81       cam_changer mc2        media_changer generic

DEVICE NODE
ID  LBdevT  LCdevT  CBdevT  CCdevT  BFlags  CFlags  Class  Suffix  L3B  L3
-----
0   0       56008c0 0       13003b3 0x0     0x861  0x0   (null) (null) (null)

COMPONENT INCONSISTENCY
-----
Component should not have an entry in the cluster database but it does.
```

3.14.3 The `hwmgr -scan scsi` Command Does Not Work Clusterwide

The `hwmgr -scan scsi` command `-cluster` option does not work clusterwide. When entered on a cluster member, the `hwmgr -scan scsi -cluster` command updates the device databases for only that member, not for all cluster members.

To perform a clusterwide scan, enter the following command on *each* cluster member when you need to update the member device databases clusterwide:

```
# hwmgr -scan comp -cat scsi_bus
```

You usually use this command when you add a new disk to a cluster (see Section 3.14.4).

3.14.4 Adding a Disk to a Running Cluster

When you add a new disk to a running cluster (for example, when you replace a failed disk), the cluster may not properly identify or configure the disk. To ensure that all cluster members properly recognize a new disk, follow these steps:

1. For all disk models, enter the following command on each member to scan SCSI buses clusterwide and configure any new devices.

```
# hwmgr -scan comp -cat scsi_bus
```

Allow a minute or two for the scans to complete.

2. If the disk that you are adding is a RZ26, RZ28, RZ29, or RZ1CB-CA model, enter the following command on each cluster member after you install the disk:

```
# /usr/sbin/clu_disk_install
```

Note

This command may take several minutes to complete if the cluster has a large number of storage devices.

3.14.5 A Media Changer Device Can Be Managed Only from Its DRD Server

To manipulate a media changer device in a cluster, you must use the `mcutil` utility on the cluster member that is the device request dispatcher (DRD) server for the device. Use the `drdmgr -a SERVER device-name` command, where *device-name* is the name of the media changer.

3.14.6 The `sysman -cli` Command Requires the `-focus` Option to Remove Member-Specific NFS Configuration Values

When you use the `sysman -cli` command in a cluster to remove member-specific NFS configuration values from a member of the cluster, you must specify the `-focus` option. For example:

```
# sysman -cli -focus member-one -set values -comp nfsconfig \  
-group client -attr cluster_override=remove -key1 member-one
```

This does not apply to the SysMan graphical user interface (GUI) or to the use of the `sysman -cli` command on a single system.

3.14.7 Running Process Accounting on Large Clusters Can Exhaust Member Process Quotas

If process accounting is enabled on large clusters (six to eight members), cluster members may start swapping heavily and eventually exhaust their process quotas. A `ps` command on such a member will show tens of thousands of `icssvr_daemon_from_pool` processes.

If you see this situation developing in a cluster that is running process accounting, use the `accton` command with no parameters to disable accounting.

3.14.8 Login Refused on Member in Multiuser Mode After Performing shutdown now Command on Lead Cluster Member

After you shut down and reboot the cluster, the `/etc/nologin_hostname` files may remain. Remove the `/etc/nologin` files to allow login access to members in multiuser state.

3.14.9 Clearing Persistent Reservations on Disks Behind HSZ80 and HSG80 Controllers

TruCluster Server cluster members use persistent reservations to coordinate access to shared storage. These persistent reservations control access to devices and logical volumes, and are used to erect a barrier against any system that is not a member of the current cluster.

The Tru64 UNIX base operating system does not know how to handle these persistent reservations. Attempts by the base operating system to access a disk that has a persistent reservation will return I/O error messages. All I/O attempts will fail.

If you boot (or reinstall) the base operating system on a system that attempts to access disks that were previously used by a cluster behind an HSZ80 or HSG80 controller, the base operating system might not be able to see the disks. This is most likely because the disk participated previously in a cluster, and one of the cluster members set a reservation on the disk to prohibit any other hosts (or non-cluster members) on this shared bus from accessing the disk. The problem is that the reservation was not cleared when the cluster member or the cluster was shut down.

To remove the reservations, use the `/usr/sbin/cleanPR clean` command. The `cleanPR` script finds and clears all persistent reservations from the attached HSZ80 and HSG80 devices.

Caution

Do not run the `cleanPR` script on a system that is in a cluster. You do not want to remove reservations that the running cluster is using to actively control disk access. Doing so can result in data corruption.

There are two scenarios to consider. The first is where you can boot a disk containing Tru64 UNIX Version 5.1 and run the `cleanPR` command. The second is where you are trying to install Tru64 UNIX Version 5.1 on a disk that is behind an HSx80 controller, has a persistent reservation on it, and there is no installed version of the base operating system available.

For the first scenario, if you encounter persistent reservations when attempting to access storage that was previously used in a cluster, boot Tru64 UNIX Version 5.1 and enter the `/usr/sbin/cleanPR clean` command to clear the reservations.

The second scenario is more complex because you do not have a running version of the operating system from which you can run `cleanP`. However, the `cleanPR` command is on the Tru64 UNIX Version 5.1 CD-ROM. Start the installation from CD-ROM. Get to the single-user shell. Enter the `/usr/sbin/cleanPR clean` command to clear the reservations.

3.15 SysMan Menu Note

The following note discusses a known problem that you may encounter when you use SysMan Menu in a cluster.

3.15.1 Nonroot Users Cannot Run System Management Applications

Most system management applications require `root` privileges to make configuration changes. Nonroot users are permitted to run system management applications only to view the current configuration. They are prevented from changing the configuration.

In a cluster, the system management applications use the remote shell command (`rsh`) to execute commands at a remote host. Part of the `rsh` command processing includes verifying access in the remote user's `$HOME/.rhosts` file in their home directory. For this reason, a nonroot user, without a home directory, who runs a system management application might encounter a core dump. Users can avoid these problems by ensuring that they have home directories set up before attempting to use the system management applications.

An additional problem exists where certain SysMan Menu tasks can be run only by the `root` user, regardless of whether a nonroot user's home directory has been established.

Each task in the SysMan Menu is associated with an action name. A task's action name is the same as its accelerator (displayed by the `sysman -list` command). These action names are associated with privileges by the Configure Division of Privilege (DOP) application in the SysMan Menu. Nonroot users can be granted the privilege to run specific actions.

Several SysMan Menu tasks do not have a required privilege associated with their actions. Such tasks can only be performed when a user is logged in as `root`. These tasks are:

- View hardware hierarchy
- View cluster
- View device information
- Manage Cluster File System
- Manage DRD Storage
- Cluster Alias Manager

3.16 SysMan Station Notes

The following notes discuss known problems that you may encounter when using SysMan Station in a cluster.

3.16.1 Unable to Expand Host Object in a Cluster

The SysMan Station client may occasionally encounter a Java class exception error when you attempt to expand a Host object.

If you encounter this error, click on a different view and then reselect the Hardware view and retry the expand operation. If the problem persists, restart the SysMan Station client and retry the expand operation.

3.16.2 SysMan Station Might Display Cluster Status Incorrectly

The SysMan Station relies on events that the Event Management (EVM) subsystem generates to monitor and display cluster status. In the following situations, the SysMan Station may reflect the state of the system incorrectly:

- The Filesystems light in the Monitor window may indicate a warning state (yellow) after all file system objects have returned to a normal state. This situation may occur after a new member has been added to

the cluster. To clear this warning, restart the SysMan Station daemon (`smsd`) on the affected cluster members by following these steps:

1. Close all open SysMan Station sessions.
2. Enter the following command:

```
# /sbin/init.d/smsd restart
```

- After a cluster member has booted, the Network light in the Monitor window may indicate a warning state (yellow) when no network errors exist. This condition is caused by network events that are generated during the boot sequence. To clear this warning, follow these steps:
 1. Click on the Network light in the Monitor window to display the Network Event window.
 2. Click on the Clear Events button.
- If the cluster application availability daemon (`caad`) fails to start on a cluster member, the SysMan Station will not correctly display the state of CAA objects. For example, this situation can happen when the CAA license is not loaded on all the cluster members. To obtain accurate information on CAA applications from the SysMan Station, follow these steps:
 1. Start the `caad` daemon on the affected cluster members using the following command:

```
# /usr/sbin/caad
```
 2. Restart the SysMan Station daemon (`smsd`) using the following command:

```
# /sbin/init.d/smsd restart
```

Additionally, the Filesystem Attention group in the SysMan Station Monitor window does not properly update on all cluster members for suboptimal file system states. If a file system becomes suboptimally configured, the Filesystem Attention group on only one cluster member will reflect the new state properly. SysMan Station clients that are connected to cluster members will not reflect this change in the Monitor window. However, the Physical Filesystem View on all cluster members will properly display this state information.

To correct this problem, stop and restart the SysMan Station daemon (`smsd`) on each cluster member where the Filesystems Attention Group is not reflecting the suboptimal state, as follows:

1. Close all open SysMan Station sessions.
2. Restart the SysMan Station daemon using the following command:

```
# /sbin/init.d/smsd restart
```

3.16.3 SysMan Station Might Display New Hardware Objects Incorrectly

If a new disk device is added or an existing disk device is replaced in a running cluster, the SysMan Station's Hardware View may display the new or modified disk object incorrectly. The disk object may be positioned incorrectly in the hardware hierarchy; for example, the disk may be drawn as a child of the host object instead of as a child of a SCSI bus.

To correct the view, restart the SysMan Station daemon (`smsd`) on each cluster member by performing the following steps on all affected members:

1. Close all open SysMan Station sessions.
2. Enter the following command:

```
# /sbin/init.d/smsd restart
```

3.16.4 Properties Might Not Be Displayed for Selected Objects

Properties may not be displayed for selected objects. The `Properties` dialog box may flash quickly on the screen or may not be displayed at all.

To work around this problem, continue to try and display properties in the current SysMan Station client or exit the SysMan Station client and start a new SysMan Station session.

3.16.5 Some SysMan Station Applications Display Wrong Target Member Name

When the following applications are launched from the SysMan Station, their title bars incorrectly display the name of the cluster member on which the SysMan Station client is running instead of the cluster member that is the target of the application's actions:

- Security Auditing Configuration
- Network Configuration Applications
- NFS Configuration Applications
- NTP Configuration Applications
- PPP Configuration Applications

The application is directed to the correct cluster member; only the name in the title bar is incorrect.

3.17 Documentation Notes

The following note corrects an error in the TruCluster Server Version 5.1 documentation:

3.17.1 Software Installation: Incorrect No WLS and Disk Space Example

The note in step 2 (the setup stage) of the rolling upgrade procedure in Chapter 7 of the *Software Installation* manual is incorrect. The note should read:

————— **Note: No WLS and Disk Space** —————

Additional space is required in the `cluster_root` domain for backing up member files on clusters without Worldwide Language Support (WLS). If no space is available, the following message is displayed:

```
*** Error ***
There is no space available in the root (/), /usr, or /var
file systems to back up member '???' member-specific files.
Increase the available disk space on one of these file systems
and rerun this stage of the upgrade.
```

The minimum required available space in the `cluster_root` domain must be greater than the sum of all of the member directories in the root (`/`), `/usr`, or `/var` file systems.

To view the used and available space in the `cluster_root` domain, enter the following command:

```
# df /
```

For example:

```
# df /
Filesystem      512-blocks  Used Available Capacity Mounted on
cluster_root#root  524288 175710   330512    35%  /
```

To calculate the minimum required value, enter the following command:

```
# ksh 'du -s {,/usr,/var}/cluster/members/member?*/' | \
awk '{minimum+=$1}; END{print minimum}'
```

For example:

```
# ksh 'du -s {,/usr,/var}/cluster/members/member?*/' | \
awk '{minimum+=$1}; END{print minimum}'
679030
```

The algorithm used by the setup stage adds the minimum required value to the amount of disk space used in the `cluster_root` domain. This sum must be less than the amount

available. Using the sample values, the `cluster_root` domain needs at least 524228 more blocks of disk space: $(679030 + 175710) - 330512 = 524228$. Use the `addvol` command to add additional volumes to the `cluster_root` domain.

This chapter provides information that was not included in the TruCluster Server Version 5.1 manuals.

4.1 Managing Third-party Storage

When a cluster member loses quorum, all of its I/O is suspended, and the remaining members erect I/O barriers against nodes that have been removed from the cluster. This I/O barrier operation inhibits non-cluster members from performing I/O with shared storage devices.

The method that is used to create the I/O barrier depends on the types of storage devices that the cluster members share. In certain cases, a Task Management function called a `Target_Reset` is sent to stop all I/O to and from the former member. This Task Management function is used in either of the following situations:

- The shared SCSI device does not support the SCSI Persistent Reserve command set and uses the Fibre Channel interconnect.
- The shared SCSI device does not support the SCSI Persistent Reserve command set, uses the SCSI Parallel interconnect, is a multiported device, and does not propagate the SCSI `Target_Reset` signal.

In either of these situations, there is a delay between the `Target_Reset` and the clearing of all I/O pending between the device and the former member. The length of this interval depends on the device and the cluster configuration. During this interval, some I/O with the former member might still occur. This I/O, sent after the `Target_Reset`, completes in a normal way without interference from other nodes.

During an interval configurable with the `drd_target_reset_wait` kernel attribute, DRD suspends all new I/O to the shared device. This period allows time to clear those devices of the pending I/O that originated with the former member and were sent to the device after it received the `Target_Reset`. After this interval passes, the I/O barrier is complete.

The default value for `drd_target_reset_wait` is 30 seconds, which should be sufficient. However, if you have doubts because of third-party devices in your cluster, contact the device manufacturer and ask for the

specifications on how long it takes their device to clear I/O after the receipt of a `Target_Reset`.

You can set `drd_target_reset_wait` at boot time and run time.

For more information about quorum loss and system partitioning, see the chapter on the connection manager in the *TruCluster Server Technical Overview*.

Index

A

- accounting
 - can exhaust process quotas in large clusters, 3-18
- addvol
 - member root domain, 2-2
- AlphaServer support, 1-1t
- AutoFS support, 1-1t

B

- Bad configuration message
 - from `clu_delete_member`, 3-6
- `bcheckrc` command, 3-5
- booting
 - ATTENTION message during, 3-3
 - CNX message during, 3-3
 - `ict_mct` messages during, 3-4
 - Memory Channel messages on two-member clusters, 3-4
 - retrying to lock boot partition during, 3-3

C

- CAA
 - action scripts PATH considerations, 3-10
 - `caa_register -u` command does not correctly update a nonapplication resource's state, 3-10
 - `caa_stat` tabular output corrupted by very long resource names, 3-12

- changes to `HOSTING_MEMBERS` list can leave applications running on disallowed members, 3-10
- erroneously lets you update a running application's required resources, 3-11
- event messages are malformed when displayed by the EVM viewer, 3-11
- fails to correctly manage applications on a member during a rolling upgrade, 3-11
- problems with `caa_relocate -f` when multiple interdependent applications are specified, 3-9
- `caa_register` command
 - erroneously lets you update a running application's required resources, 3-11
 - `-u` option does not correctly update a nonapplication resource's state, 3-10
- `caa_stat`
 - tabular output corrupted by very long resource names, 3-12
- CDFS
 - supported for read-only access, 2-1
- CFS
 - concurrent direct I/O support, 1-1t
 - relocation failure, 3-7
 - restrictions on, 2-1
- `cfsmgr` command
 - dumps core when `-p` or `-m` option is used with `-F raw` option, 3-8

- fails to relocate file systems
 - containing applications that wire memory, 3-7
- cleanPR command, 3-18
- clu_add_member command, 3-2
 - if kernel build for newly added member fails on first boot, 3-3
- clu_check_config command
 - ATTENTION message from, 3-3
- clu_create command, 3-2
- clu_delete_member command
 - Bad configuration message from, 3-6
 - cannot delete member with a boot disk with a bad disk label, 3-6
 - does not adjust expected votes when deleting a member with inaccessible boot disk, 3-6
 - NTP and mail messages from, 3-7
- cluamgr command
 - can enable or disable only one interface at a time, 3-13
- cluster alias
 - aliasd daemon sets up duplicate subnets, 3-12
 - cluamgr command, 3-13
- cluster configuration check
 - ATTENTION message from, 3-3

D

- devices
 - managing third-party storage, 4-1
- disabling Internet services, 3-15
- disk
 - replacing a failed, 3-17
- disks
 - persistent reservations, 3-18
- documentation
 - errors in, 3-23
- drd_target_reset_wait attribute, 4-1

E

- /etc/cfgmgr.auth file
 - adding first member to, 3-2

F

- features, 1-1
- FFM file system
 - not supported, 2-2
- floppy disks
 - and UFS read-only file system, 2-1

H

- HSG80 controller
 - persistent reservations, 3-18
- HSZ80 controller
 - persistent reservations, 3-18
- hwmgr command
 - cluster option does not work clusterwide, 3-16
 - show comp option may report an inconsistency error when creating clusterwide names for SCSI devices, 3-16

I

- I/O barrier
 - drd_target_reset_wait attribute, 4-1
 - managing third-party storage, 4-1
- icssvr_daemon_from_pool processes, 3-18
- inetd
 - disabling Internet services, 3-15
- init -s command, 3-5
- installation of cluster software
 - do not use SysMan Menu to perform, 3-1
- Internet services
 - disabling, 3-15

K

- kernel build
 - if fails for newly added member on first boot, 3-3
- kernel configuration management database
 - adding first member to, 3-2
- kloadsrv command, 3-5

L

- lmf reset command, 3-5
- LSM
 - voldisk list command does not display all non-LSM disks within a cluster, 3-9

M

- mail
 - messages from clu_delete_member command, 3-7
- mcutil command
 - must be used on DRD server of media changer, 3-17
- media changer devices
 - can be managed only from their DRD servers, 3-17
- member root domain
 - addvol command, 2-2
- Memory Channel
 - errors on two-member clusters, 3-4
 - errors when booting a member into a cluster, 3-4
- messages
 - ATTENTION from boottime cluster check, 3-3
 - Bad configuration message from clu_delete_member command, 3-6

- CNX invalid configuration for seq disk, 3-3
- ict_mct, 3-4
- mail messages from
 - clu_delete_member command, 3-7
- Memory Channel (rm), 3-4
- NTP messages from
 - clu_delete_member command, 3-7
- Server Relocation Failed from
 - cfsmgr command, 3-8
 - trying to lock boot partition, 3-3
- MFS
 - not supported, 2-1
- mlock system call
 - CFS relocation problems, 3-7

N

- named pipes
 - not supported, 2-2
- network interfaces
 - enabling or disabling with cluamgr, 3-13
- NFS client file system
 - failover not supported, 2-1
- NFS file system
 - unmounted unexpectedly, 3-7
- NTP
 - messages from clu_delete_member command, 3-7

P

- partitioning
 - drd_target_reset_wait attribute, 4-1
- persistent reservations, 3-18
- plock system call
 - CFS relocation problems, 3-7
- /proc file system

not supported, 2-1
process accounting
 can exhaust process quotas in large
 clusters, 3-18

R

rolling upgrade
 clu_upgrade preinstall ignores user
 response, 3-14
 do not run setld after clu_upgrade
 setup, 3-14
 login refused in multiuser mode
 after shutdown now, 3-18
 messages you can ignore, 3-14
 run dop only on rolled members,
 3-15
 updating of WLS from CD-ROM,
 3-14
 version not updated in /etc/motd,
 3-15

S

Server Relocation Failed message
 from cfsmgr command, 3-8
setld command
 do not use between clu_upgrade
 setup and rolling the first
 member, 3-14
shutdown -s command
 do not use, 3-5
shutdown command
 for single member unmounts
 inactive NFS file systems, 3-7
single-user mode
 first halt member, then boot to, 3-5
storage
 managing third-party devices, 4-1
sysman -cli command
 using -focus to remove
 member-specific NFS
 configuration values, 3-17

SysMan Menu
 do not use Installation branch of
 Software menu of, 3-1
 nonroot users cannot perform
 management tasks, 3-19

SysMan Station

 applications display wrong target
 member name, 3-22
 Filesystem Attention group may not
 be updated for nonoptimal file
 system states, 3-21
 may display cluster status
 incorrectly, 3-20
 may display new hardware objects
 incorrectly, 3-22
 properties may not display for
 selected objects, 3-22
 shows CAA application resources in
 UNKNOWN state as having an
 error, 3-12
 unable to expand Host object in a
 cluster, 3-20

U

UFS
 and floppy disks, 2-1
 supported for read-only access, 2-1

V

voldisk command
 list command does not display all
 non-LSM disks within a cluster,
 3-9

W

wired memory
 CFS relocation problems, 3-7

How to Order Tru64 UNIX Documentation

To order Tru64 UNIX documentation in the United States and Canada, call **800-344-4825**. In other countries, contact your local Compaq subsidiary.

If you have access to Compaq's intranet, you can place an order at the following Web site:

<http://asmorder.nqo.dec.com/>

If you need help deciding which documentation best meets your needs, see the Tru64 UNIX *Documentation Overview*, which describes the structure and organization of the Tru64 UNIX documentation and provides brief overviews of each document.

The following table provides the order numbers for the Tru64 UNIX operating system documentation kits. For additional information about ordering this and related documentation, see the *Documentation Overview* or contact Compaq.

Name	Order Number
Tru64 UNIX Documentation CD-ROM	QA-6ADAA-G8
Tru64 UNIX Documentation Kit	QA-6ADAA-GZ
End User Documentation Kit	QA-6ADAB-GZ
Startup Documentation Kit	QA-6ADAC-GZ
General User Documentation Kit	QA-6ADAD-GZ
System and Network Management Documentation Kit	QA-6ADAE-GZ
Developer's Documentation Kit	QA-6ADAF-GZ
Reference Pages Documentation Kit	QA-6ADAG-GZ
TruCluster Server Documentation Kit	QA-6BRAA-GZ

Reader's Comments

TruCluster Server

Release Notes

AA-RHH1C-TE

Compaq welcomes your comments and suggestions on this manual. Your input will help us to write documentation that meets your needs. Please send your suggestions using one of the following methods:

- This postage-paid form
- Internet electronic mail: readers_comment@zk3.dec.com
- Fax: (603) 884-0120, Attn: UBPG Publications, ZKO3-3/Y32

If you are not using this form, please be sure you include the name of the document, the page number, and the product name and version.

Please rate this manual:

	Excellent	Good	Fair	Poor
Accuracy (software works as manual says)	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Clarity (easy to understand)	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Organization (structure of subject matter)	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Figures (useful)	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Examples (useful)	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Index (ability to find topic)	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Usability (ability to access information quickly)	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Please list errors you have found in this manual:

Page	Description
_____	_____
_____	_____
_____	_____
_____	_____

Additional comments or suggestions to improve this manual:

What version of the software described by this manual are you using? _____

Name, title, department _____

Mailing address _____


Electronic mail _____

Telephone _____


Date _____

----- Do Not Cut or Tear - Fold Here and Tape -----

COMPAQ



NO POSTAGE
NECESSARY IF
MAILED IN THE
UNITED STATES



BUSINESS REPLY MAIL
FIRST CLASS MAIL PERMIT NO. 33 MAYNARD MA

POSTAGE WILL BE PAID BY ADDRESSEE

COMPAQ COMPUTER CORPORATION
UBPG PUBLICATIONS MANAGER
ZKO3-3/Y32
110 SPIT BROOK RD
NASHUA NH 03062-2698



----- Do Not Cut or Tear - Fold Here -----

Cut on This Line