

TruCluster Software Products

Release Notes

Part Number: AA-R0JAD-TE

April 1999

Product Version: TruCluster Production Server
Software Version 1.6, TruCluster
Available Server Software Version 1.6,
and TruCluster Memory Channel
Software Version 1.6

Operating System and Version: Tru64 UNIX (formerly DIGITAL
UNIX) Version 4.0F

This manual provides important information about the TruCluster software products.

© Digital Equipment Corporation 1998, 1999
All rights reserved.

COMPAQ, the Compaq logo, and the Digital logo are registered in the U.S. Patent and Trademark Office. The following are trademarks of Digital Equipment Corporation: ALL-IN-1, Alpha AXP, AlphaGeneration, AlphaServer, AltaVista, ATMworks, AXP, Bookreader, CDA, DDIS, DEC, DEC Ada, DECevent, DEC Fortran, DEC FUSE, DECnet, DECstation, DECsystem, DECterm, DECUS, DECwindows, DTIF, Massbus, MicroVAX, OpenVMS, POLYCENTER, PrintServer, Q-bus, StorageWorks, Tru64, TruCluster, TURBOchannel, ULTRIX, ULTRIX Mail Connection, ULTRIX Worksystem Software, UNIBUS, VAX, VAXstation, VMS, and XUI. Other product names mentioned herein may be the trademarks of their respective companies.

Prestoserve is a trademark of Legato Systems, Inc.; the trademark and software are licensed to Digital Equipment Corporation by Legato Systems, Inc. NFS is a registered trademark of Sun Microsystems, Inc. Open Software Foundation, OSF, OSF/1, OSF/Motif, and Motif are trademarks of the Open Software Foundation, Inc. UNIX is a registered trademark and The Open Group is a trademark of The Open Group in the US and other countries. Oracle is a registered trademark of Oracle Corporation. Netscape and Netscape Navigator are registered trademarks of Netscape Communications Corporation in the United States and other countries. Baan is a trademark of Baan Development B.V. INFORMIX is a registered trademark of Informix Software, Inc. Tuxedo is a registered trademark of Novell, Inc.

Restricted Rights: Use, duplication, or disclosure by the U.S. Government is subject to restrictions as set forth in subparagraph (c) (1) (ii).

Compaq Computer Corporation makes no representations that the use of its products in the manner described in this publication will not infringe on existing or future patent rights, nor do the descriptions contained in this publication imply the granting of licenses to make, use, or sell equipment or software in accordance with the description.

Possession, use, or copying of the software described in this publication is authorized only pursuant to a valid written license from Compaq or an authorized sublicensor.

Compaq conducts its business in a manner that conserves the environment and protects the safety and health of its employees, customers, and the community.

Contents

About This Manual

1 Release Notes

1.1	New Features	1-1
1.2	Hardware and Firmware Requirements	1-4
1.2.1	Supported Member Systems	1-4
1.2.2	SCSI-2 Adapters Supported for Shared Storage	1-6
1.2.3	Supported Memory Channel Hardware	1-6
1.2.4	Fibre Channel Hardware Supported for Cluster Storage ..	1-7
1.2.5	Supported RAID Controllers	1-8
1.2.6	Supported SCSI Signal Converters	1-9
1.2.7	Supported Network Adapters	1-9
1.2.7.1	Supported Ethernet and FDDI Adapters	1-9
1.2.7.2	Supported ATM Adapters	1-10
1.2.8	Supported Disks	1-10
1.2.9	Supported Tape Devices	1-11
1.3	List of Subsumed Patches	1-12
1.4	KZPSA without Devices Causes Machine Check During Shutdown	1-12
1.5	Do Not Use ifconfig or nfs_ifconfig to Configure Multiple Addresses for an NFS Service	1-13
1.6	The use of ifconfig down is not a Replacement for Cable Pulls	1-14
1.7	Tuning NFS Retransmission Timeout	1-14
1.8	The cpq_mibs Daemon May Generate Binary Error Log Entries in an ASE Environment	1-14
1.9	Memory Channel Notes and Restrictions	1-15
1.10	Fibre Channel Requirements and Restrictions	1-16
1.11	UltraSCSI Hub Restrictions	1-18
1.12	Enhanced Security in an Available Server or Production Server Configuration	1-18
1.12.1	Issues with Distributed Logins, Enhanced Security, and NIS	1-18
1.12.2	An ASE Solution Using NIS and a Highly Available NFS Service	1-19

1.12.3	Configuring Distributed Enhanced Security with the TruCluster Software Products	1-21
1.12.3.1	Setting Up the Initial (NIS Master) System	1-21
1.12.3.2	Setting Up Additional Member Systems	1-25
1.12.3.3	Completing the Distributed Accounts Set Up on the NIS Master	1-26
1.12.3.4	Completing the Distributed Accounts Set Up on the Additional Member Systems	1-28
1.12.4	Audit in an ASE Configuration	1-28
1.12.5	Available Server and Production Server NFS File Lock Configuration	1-28
1.12.5.1	Purpose	1-29
1.12.5.2	Procedure	1-29
1.12.6	Administration Notes	1-29

Index

Examples

1-1	Preventing the cpq_mibs Daemon From Starting	1-15
-----	--	------

Tables

1-1	New Features in TruCluster Software Products Version 1.6 ..	1-1
1-2	Supported Member Systems	1-5
1-3	SCSI-2 Adapters Supported for Shared Storage	1-6
1-4	Supported Memory Channel Hardware	1-7
1-5	Fibre Channel Hardware Supported for Cluster Storage	1-8
1-6	Supported RAID Controllers	1-8
1-7	Supported Disk Devices	1-10
1-8	Supported Tape Devices	1-12
1-9	AlphaServer Systems Supported for Fibre Channel	1-16

About This Manual

This manual provides important information about the TruCluster™ software products.

Audience

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

Organization

This manual contains a chapter of release notes and an index.

Related Documents

Users of the TruCluster Production Server Software can consult the following TruCluster Software Products manuals for assistance in cluster hardware configuration, installation, administration, and programming tasks:

- *Hardware Configuration*—Describes how to set up the systems that are to become cluster members, and how to configure cluster shared storage.
- *Software Installation*—Describes how to install the TruCluster Software on the systems that are to participate in the cluster.
- *Administration*—Describes cluster-specific administration tasks, such as those required to set up an available server environment (ASE) within a cluster. It also shows how to configure, start, and manage distributed raw disk (DRD) services and other available services.
- TruCluster Production Server Software *Application Programming Interfaces*—Describes the application programming interfaces (APIs) provided by the distributed lock manager (DLM) and cluster information services.
- TruCluster Production Server Software *Memory Channel Application Programming Interfaces*—Describes the APIs used to program the features of the Memory Channel hardware.

In addition to these release notes, users of the TruCluster Memory Channel Software should consult only the TruCluster Software Products *Hardware*

Configuration and TruCluster Production Server Software *Memory Channel Application Programming Interfaces* manuals.

Users of the TruCluster Available Server Software should consult only the TruCluster Software Products *Hardware Configuration*, *Software Installation*, and *Administration* manuals for assistance in configuring, installing, and managing an available server environment (ASE).

Online Documentation

Each book in the TruCluster Software documentation set is shipped as a set of Hypertext Markup Language (HTML) and graphics files in the /TCR/doc/html directory on the Associated Products Volume 2 CD-ROM. You can use the Netscape® Navigator™ World Wide Web browsing program to display these books.

If the Tru64 UNIX installation program detects graphics capabilities on your system, it automatically installs Netscape Navigator. You can then invoke Netscape from an icon on the Common Desktop Environment (CDE) front panel or directly from the command line. Detailed help for Netscape is available through the help menus.

To access the TruCluster Software documentation from the viewer, click on the Open icon in the Netscape main window and enter the following file location in the Open Location: text entry field:

```
file:/mountpoint/TCR/doc/html/BOOKSHELF.HTM
```

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 03062-9987

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

This manual uses the following typographical conventions:

#	A number sign represents the superuser prompt.
% cat	Boldface type in interactive examples indicates typed user input.
<i>file</i>	Italic (slanted) type indicates variable values, placeholders, and function argument names.
⋮	A vertical ellipsis indicates that a portion of an example that would normally be present is not shown.
cat(1)	A cross-reference to a reference page includes the appropriate section number in parentheses. For example, <code>cat(1)</code> indicates that you can find information on the <code>cat</code> command in Section 1 of the reference pages.
PS	Abbreviation for the TruCluster Production Server Software.

AS Abbreviation for the TruCluster Available Server Software.

MC Abbreviation for the TruCluster Memory Channel Software.

1

Release Notes

These release notes provide important information about Version 1.6 of the TruCluster software products.

1.1 New Features

TruCluster Software Products Version 1.6 introduces the features listed in Table 1–1. The table indicates which features are provided by TruCluster Production Server Software (PS), TruCluster Memory Channel Software (MC), and TruCluster Available Server Software (AS). The abbreviation "N/A" indicates that the listed feature is not applicable to a particular product.

Table 1–1: New Features in TruCluster Software Products Version 1.6

Feature	PS	MC	AS	For more information, see:
Ability to modify and retain DRD device special file permissions, owner, and group settings.	Yes	N/A	N/A	<i>Administration</i>
Adds support to allow clients of an ASE NFS service to mount the exported filesystem using TCP as the network transport.	Yes	N/A	Yes	<i>Tru64 UNIX Network Administration</i>
Provides the ability for the ASE manager to add services based on the contents of service configuration files.	Yes	N/A	Yes	<i>Administration</i>
Provides a mechanism to modify the internal, nonuser-defined ASE default action script timeout values.	Yes	N/A	Yes	<i>Administration</i>
Provides the ability to monitor up to eight network interfaces.	Yes	N/A	Yes	<i>Administration</i>
Adds the ASE_IGNORE_BCL configuration variable. If the disks used for LSM Block Change Logging (BCL) are not available, the services will still start if the ASE_IGNORE_BCL configuration variable is set to on.	Yes	N/A	Yes	<i>Administration</i>

Table 1–1: New Features in TruCluster Software Products Version 1.6 (cont.)

Feature	PS	MC	AS	For more information, see:
Adds support for the <code>cnxmibd</code> daemon, an extensible SNMP (eSNMP) subagent that implements the Compaq Common Cluster MIB.	Yes	Yes	Yes	<code>cnxmibd</code> reference page
Supports enhanced security with distributed authentication.	Yes	N/A	Yes	Tru64 UNIX <i>Security</i>
Provides a procedure and the tools necessary to recover from a panicking ASE.	Yes	N/A	Yes	<i>Administration</i>
Provides the ability to enable use of the <code>fuser -k</code> command to kill all processes on a file system mount point before the unmount is attempted on a service stop request.	Yes	N/A	Yes	<i>Administration</i> and <code>fuser</code> reference page
Provides the <code>ASE_ADVFS_D_IGNORE_DISKS</code> and <code>ASE_SNMPD_IGNORE_DISKS</code> configuration variables. When the variable is set to <code>yes</code> , it prevents the <code>advfsd</code> or <code>os_mibs</code> daemons from requesting that the kernel access device special files for devices reserved by an ASE service running on another system.	Yes	N/A	Yes	<i>Administration</i>
Provides supported example scripts, which can be modified as necessary to set up highly available NIS master or print server services.	Yes	N/A	Yes	<i>Administration</i> , <code>/var/ase \</code> <code>/examples_supported \</code> <code>/README</code>
Provides unsupported example scripts to start and stop:	Yes	N/A	Yes	<code>/var/ase \</code> <code>/examples_unsupported \</code> <code>/README</code>
<ul style="list-style-type: none"> • An Oracle® database • An INFORMIX® database • The AltaVista search intranet daemons • A Netscape® Web server • A Netscape News server • A SAMBA PC fileserver • An Apache Web server • A Baan™ environment • A Tuxedo® environment^a 				

Table 1–1: New Features in TruCluster Software Products Version 1.6 (cont.)

Feature	PS	MC	AS	For more information, see:
Includes the introduction of the Reliable Datagram (RDG) messaging facility. RDG provides applications developed to its defined application programming interface (API) with higher bandwidth and lower latencies than are obtainable with IP-based protocols. Future releases from database partners are planned to include support for RDG.	Yes	N/A	N/A	N/A ^b
Adds ASE manager support for the mount command dual option, which allows an AdvFS split mirror to be mounted.	Yes	N/A	Yes	<i>Administration</i>
Support for Memory Channel 2 hardware operating at 512-MB Memory Channel address space. Support includes Memory Channel fiber optics converters used to increase the distance between ASEs or cluster members.	Yes	Yes	N/A	Section 1.2.3 and <i>Hardware Configuration</i>
Support for Fibre Channel for cluster storage	Yes	N/A	Yes	Section 1.2.4, Section 1.2.5, Section 1.10, and <i>Hardware Configuration</i>
Provides support for mixed Memory Channel interconnects operating at 128-MB Memory Channel address space. Also provide the ability to initiate the use of 512-MB Memory Channel address space after a cluster has been rolled from MC1 or MC1.5 to MC2 hardware through the use of the Memory Channel kernel subsystem <code>rm_use_512</code> entry.	Yes	Yes	N/A	Section 1.9, <i>Hardware Configuration</i> , and <i>Administration</i>
Support redundant array of independent network adapters (NetRAIN) devices as the cluster interconnect for TruCluster Available Server Software and as the monitored client network interfaces for both TruCluster Available Server Software and TruCluster Production Server Software.	Yes	N/A	Yes	Tru64 UNIX Version 4.0F <i>Release Notes</i> and Tru64 UNIX <i>Network Administration</i>
Support for the KZPBA-CB UltraSCSI adapter. ^c	Yes	N/A	Yes	Section 1.2.2 and <i>Hardware Configuration</i>
Support for the RAID Array 3000 (RA3000) storage subsystem. The RA3000 uses the HSZ22 controller(s). ^c	Yes	N/A	Yes	Section 1.2.5 and <i>Hardware Configuration</i>

Table 1–1: New Features in TruCluster Software Products Version 1.6 (cont.)

Feature	PS	MC	AS	For more information, see:
Support for the DS-DWZZH-03 and DS-DWZZH-05 UltraSCSI hubs. ^c	Yes	N/A	Yes	Section 1.2.6 and <i>Hardware Configuration</i>
Adds support for additional tape devices on a shared SCSI bus with server failover for the NetWorker product.	Yes	N/A	Yes	Section 1.2.9, Table 1–8, <i>Hardware Configuration, and Administration</i>

^aThe Tuxedo scripts are written in Perl and require a perl5 binary in the /usr/ucb/perl directory

^bThe documentation of the API for use by ISVs is being considered for future releases of TruCluster software.

^cSupport for these devices was first announced for DIGITAL UNIX Version 4.0D and TruCluster Software Products Version 1.5. This is the first new release of the TruCluster Software products that supports these devices. The information provided in the technical updates has been incorporated into the *Hardware Configuration* manual.

You cannot take advantage of Version 1.6 functionality until the entire cluster has been upgraded to Version 1.6. For example, you cannot enable clients to mount an ASE NFS service filesystem using TCP as the network transport until the upgrade is complete and you have enabled TruCluster Version 1.6 functionality. Please read the TruCluster Software Products *Software Installation* manual before installing any of the TruCluster software products.

1.2 Hardware and Firmware Requirements

You must set up the TruCluster hardware configuration before you install any TruCluster software product. This section lists the hardware supported by (and, in some cases, required by) Version 1.6 of each TruCluster software product. To obtain more detailed information on the role any hardware component plays in a cluster or available server environment (ASE), as well as configuration assistance, see the TruCluster Software Products *Hardware Configuration* manual. That manual also provides restrictions that apply to the use of a specific piece of hardware and procedures for verifying the firmware revision level of individual hardware components.

1.2.1 Supported Member Systems

Table 1–2 lists the systems supported by the TruCluster Version 1.6 software products, which SCSI adapter is supported on the system, and indicates which systems are supported by TruCluster Production Server Software (PS), TruCluster Memory Channel Software (MC), and TruCluster Available Server Software (AS). See the pertinent Software Product Description (SPD) for a detailed list of supported models in each system family. All systems must be running Tru64 UNIX Version 4.0F.

Table 1–2 also lists the minimum Alpha System Reference Manual (SRM) firmware version required for each system. For systems running Tru64

UNIX Version 4.0F, firmware is distributed on the Alpha Systems Firmware Update Version 5.4 CD-ROM.

Table 1–2 also provides a list of supported member systems, the SCSI adapters supported on the system, and the minimum SRM firmware revision needed for Tru64 UNIX and the TruCluster products.

Table 1–2: Supported Member Systems

System	SCSI or Fibre Channel Adapter	Minimum SRM Firmware for Version 4.0F	PS	MC	AS
DEC 3000	PMAZC or KZTSA	V7.0	No	No	Yes
DEC 7000 and DEC 10000	KZMSA	V5.4	No	No	Yes
AlphaServer 300	KZPSA	V7.0	No	No	Yes
AlphaServer 400	KZPSA	V7.0	No	No	Yes
AlphaServer 800	KZPSA, KZPBA-CB, KGPSA-BC	V5.4	Yes	Yes	Yes
AlphaServer 1000	KZPSA	V5.4	No	No	Yes
AlphaServer 1000A	KZPSA or KZPBA-CB	V5.4	Yes	Yes	Yes
AlphaServer 1200	KZPSA, KZPBA-CB, KGPSA-BC	V5.4	Yes	Yes	Yes
AlphaServer 2000	KZPSA or KZPBA-CB	V5.3	Yes	Yes	Yes
AlphaServer 2100	KZPSA or KZPBA-CB	V5.3	Yes	Yes	Yes
AlphaServer 2100A	KZPSA or KZPBA-CB	V5.3	Yes	Yes	Yes
AlphaServer 4000	KZPSA, KZPBA-CB, KGPSA-BC	V5.4	Yes	Yes	Yes
AlphaServer 4100	KZPSA, KZPBA-CB, KGPSA-BC	V5.4	Yes	Yes	Yes
DS20	KZPSA, KZPBA-CB, KGPSA-BC	V5.4	Yes	Yes	Yes
ES40	KZPSA, KZPBA-CB, KGPSA-BC	V5.4	Yes	Yes	Yes

Table 1–2: Supported Member Systems (cont.)

System	SCSI or Fibre Channel Adapter	Minimum SRM Firmware for Version 4.0F	PS	MC	AS
AlphaServer 8200 and 8400	KZPSA, KZPBA-CB, KGPSA-BC	V5.4	Yes	Yes	Yes
GS60 and GS140	KZPSA, KZPBA-CB, KGPSA-BC	V5.4	Yes	Yes	Yes

1.2.2 SCSI-2 Adapters Supported for Shared Storage

Table 1–3 indicates which SCSI-2 adapters are supported for shared storage by TruCluster Production Server Software Version 1.6 (PS) and TruCluster Available Server Software Version 1.6 (AS). Refer to Table 1–2 to determine if the SCSI adapter is supported on any particular system.

Table 1–3: SCSI-2 Adapters Supported for Shared Storage

SCSI-2 Adapter	Hardware Revision	Firmware Revision	PS	AS
KZPBA-CB	A01 or higher	5.57 or higher	Yes	Yes
KZPSA PCI-to-Fast-Wide Differential ^a	F01	A11 or higher	Yes	Yes
KZTSA TURBOchannel-to-Fast-Wide Differential	F01	A09 or higher	No	Yes
KZMSA XMI	F03	V5.6 or higher	No	Yes
PMAZC TURBOchannel Dual-Channel	N/A	V1.8 or higher	No	Yes

^aYou cannot update the KZPSA firmware on an AlphaServer 1000A system or 2100A system if the KZPSA adapter is located behind the PCI-to-PCI bridge.

1.2.3 Supported Memory Channel Hardware

Table 1–4 describes the Memory Channel hardware supported by TruCluster Production Server Software Version 1.6 and TruCluster Memory Channel Software Version 1.6. Memory Channel hardware is not supported by TruCluster Available Server Software.

See Section 1.9 for Memory Channel restrictions.

Table 1–4: Supported Memory Channel Hardware

MC1	MC2	Description
CCMAA-AA or CCMAA-BA	CCMAB-AA	Memory Channel PCI adapter: Used as the cluster interconnect.
CCMHA-AA	CCMHB-AA	Memory Channel hub: PC-class enclosure that is populated with linecards and used to connect Memory Channel adapters. A hub is required if you have more than two member systems. In addition to linecards, the MC2 hub may also contain fiber optic converters.
CCMLA-AA	CCMLB-AA	Linecard: Installed in a Memory Channel hub to allow connection of more than two systems.
		Copper link cable:
BC12N-10 (10 meters)	N/A	Connect a Memory Channel adapter to a hub or one Memory Channel adapter to another Memory Channel adapter (virtual hub).
N/A	BN39B-04 (4 meters) or BN39B-10 (10 meters)	Connect a Memory Channel adapter to another Memory Channel adapter (virtual hub) or a hub linecard.
N/A	BN39B-01 (1 meter)	Connect a Memory Channel adapter to a CCMFB optical converter.
N/A	CCMFB-AA	Fiber optics converter.
N/A	BN34R-10 (10 meters) or BN34R-31 (31 meters).	Fiber optics cable: Connect one optical converter to another optical converter.

1.2.4 Fibre Channel Hardware Supported for Cluster Storage

Table 1–5 lists the Fibre Channel hardware supported for cluster storage by TruCluster Production Server Software Version 1.6 (PS) and TruCluster Available Server Software Version 1.6 (AS).

Table 1–5: Fibre Channel Hardware Supported for Cluster Storage

Device	Hardware Revision	Firmware Revision
KGPSA-BC PCI-to-Fibre Channel host adapter	A01	2.20
DS-DSGGA-AA/AB 8/16 port Brocade Fibre Channel switch	A01	1.6b
HSG80 array controller	C01	V8.4

See Table 1–6 for more information on the HSG80 array controller. See Section 1.10 for Fibre Channel requirements and restrictions.

1.2.5 Supported RAID Controllers

Table 1–6 lists the RAID controllers supported by TruCluster Production Server Software Version 1.6 (PS) and TruCluster Available Server Software Version 1.6 (AS). It indicates which Hierarchical Storage Operating Firmware (HSOF) or SWCC revisions are required for each RAID controller.

Table 1–6: Supported RAID Controllers

RAID Controller	Firmware Revision	PS	AS
HSZ10- <i>Ax</i>	0306	No	Yes ^a
SWXRA-Z1 (HSZ20)	HSOF V3.0-2	Yes	Yes
HSZ40- <i>Ax</i>	HSOF V2.0 or V2.5	Yes	Yes
HSZ40- <i>Bx</i>	HSOF V2.5 or higher	Yes	Yes
HSZ40- <i>Cx</i>	HSOF V2.7 or higher	Yes	Yes
HSZ50- <i>Ax</i>	HSOF V5.0 or higher	Yes	Yes
HSZ70 ^b	HSOF V7.0	Yes	Yes
RAID array 3000 (RA3000) with HSZ22 controller	HSOF VD11s	Yes	Yes
HSZ80 array controller ^{c d}	Array Controller Software (ACS) Version 8.3Z	Yes	Yes
HSG80 array controller	Array Controller Software (ACS) Version 8.4	Yes	Yes

^aThe HSZ10 controller can be used only in conjunction with PMAZC adapters in an Available Server configuration.

^bThe HSZ70 controller can be used only in conjunction with KZPSA adapters in a Production Server cluster or an Available Server configuration.

^cThe HSZ80 controller is supported only with the AlphaServer 800, AlphaServer 1200, AlphaServer 4x00, and AlphaServer 8x00 systems with the KZPBA-CB UltraSCSI host adapter.

^dSupport for the HSZ80 was announced after the documentation for the TruCluster Software Products had been closed; therefore, there is no supporting information in the TruCluster documentation.

1.2.6 Supported SCSI Signal Converters

The following hardware revisions of SCSI signal converters are required:

- DWZZA-AA—Revision E01 or later
- DWZZA-VA—Revision F01 or later
- DWZZB-AA—Revision A01 or later
- DWZZB-VW—Revision A01 or later
- DS-DWZZH-03—Revision A01 or later
- DS-DWZZH-05—Revision E01 or later

Note

Support for the DS-DWZZH-03 and DW-DWZZH-05 was first announced for DIGITAL UNIX Version 4.0D and TruCluster Software Products Version 1.5. This is the first new release of the TruCluster Software products that supports these devices. The information provided in the technical updates has been incorporated into the TruCluster Software Products *Hardware Configuration* manual.

1.2.7 Supported Network Adapters

Ethernet, FDDI, and Asynchronous Transfer Mode (ATM) networks are supported by the TruCluster Software products. The following sections describe these adapters.

1.2.7.1 Supported Ethernet and FDDI Adapters

The following Ethernet and FDDI adapters are supported:

- DE504-BA (PCI/Quad Port Ethernet)
- DE500 (PCI/Fast Ethernet)
- DEFPA (PCI/FDDI)
- DE435 (PCI/Ethernet)
- DEFEA (EISA/FDDI)
- DE422 (EISA/Lance Ethernet)
- DE425 (EISA/Ethernet)

- DEMNA (XMI/Ethernet)
- DEMFA (XMI/FDDI)
- PMAD (TURBOchannel/Ethernet)
- DEFTA (TURBOchannel/FDDI)
- DEFZA (TURBOchannel/FDDI)
- DEGPA-SA (PCI/Gigabit Ethernet)

Note

The DEGPA-SA Gigabit Ethernet adapter is supported only with the AlphaServer 800, AlphaServer 1200, AlphaServer 4x00, and AlphaServer 8x00 systems.

1.2.7.2 Supported ATM Adapters

Asynchronous Transfer Mode (ATM) is a high-speed, connection-oriented, cell-switched technology. It meets the real-time networking requirements of multimedia applications, while providing increased bandwidth for current installations.

The DGLPB-AB ATMworks 350 155 MB/sec multimode fiber optics PCI adapter is supported.

1.2.8 Supported Disks

Table 1–7 lists the disk devices supported by TruCluster Production Server Software Version 1.6 and TruCluster Available Server Software Version 1.6 on a shared SCSI bus. It indicates the firmware levels that are known to work for each disk at the time of this manual's publication.

Table 1–7: Supported Disk Devices

Disk	Data Path	Firmware Revisions
RZ26-VA	Narrow	392 or 392A
RZ26L-VA	Narrow	442D
RZ26L-VW	Wide	442E
RZ26N-VA/RZ26N-VW	Narrow/Wide	0466, 0568, 0616, or 1103
RZ28-VA	Narrow	442C or 442D

Table 1–7: Supported Disk Devices (cont.)

Disk	Data Path	Firmware Revisions
RZ28-VW	Wide	442E
RZ28B-VA	Narrow	0006
RZ28D-VA/RZ28D-VW	Narrow/Wide	0008
RZ28L-VA/RZ28L-VW	Narrow/Wide	LYJ0
RZ28M-VA/RZ28M-VW	Narrow/Wide	0466, 0568, 0616, or 1103
RZ29B-VA/RZ29B-VW	Narrow/Wide	0014 or 0016
RZ29L-VA/RZ29L-VW	Narrow/Wide	LYJ0
RZ40-VA/RZ40-VW	Narrow/Wide	0305
RZ40L-VA/RZ40L-VW	Narrow/Wide	LYJ0
RZ1BB-VW	Wide	LYJ0, 0844, or 0845
RZ1CB-VW	Wide	LYJ0, 0847
RZ1DB-VW	Wide	LYJ0, 0307
DS-RZ1CF-VW	Wide	0371, N1H1
DS-RZ1DF-VW	Wide	0371, N1H1
DS-RZ1EF-VW	Wide	0371, N1H1
DS-RZ2CD-VW	Wide	0306
DS-RZ2DD-VW	Wide	0306
DS-RZ2ED-VW	Wide	0306
DS-RZ1CD-VW	Wide	0306
DS-RZ1BB-VW	Wide	LYJ0/0658
DS-RZ1CB-VW	Wide	LYJ0/0658
DS-RZ1DB-VW	Wide	LYJ0/0658

1.2.9 Supported Tape Devices

Table 1–8 lists the tape devices supported by TruCluster Production Server Software Version 1.6 and TruCluster Available Server Software Version 1.6 on a shared SCSI bus. It indicates the firmware levels that are known to work for each tape at the time of this manual's publication. Chapter 2 of the *Hardware Configuration* manual provides an overview of each of the supported tape devices.

Table 1–8: Supported Tape Devices

Tape Device	Supported with		Firmware Revisions	
	KZPSA	KZPBA-CB	Tape Drive	Robotics
TZ88	Yes	No	V80	N/A
TZ89	Yes	No	V80	N/A
TZ885 DLT Mini-library	Yes	No	V80	N/A
TZ887 DLT Mini-library ^a	Yes	No	V80	N/A
TL890 MiniLibrary Expansion Unit ^{a b}	Yes	Yes	N/A	N/A
TL891/892 MiniLibraries ^a	Yes	Yes	V80	3.23
TL893/896 Automated Tape Libraries ^a	Yes	No	V80	V51 ^c
TL894 Automated Tape Library ^a	Yes	No	V80	S2.2
TL895 Automated Tape Library ^a	Yes	No	V80	N2.20

^aSupport for these devices was first announced for DIGITAL UNIX Version 4.0D and TruCluster Software Products Version 1.5. This is the first new release of the TruCluster Software products that supports these devices. The information provided in the technical updates has been incorporated into the TruCluster Software Products *Hardware Configuration* manual.

^bThe TL890 contains no tape drives, but adds 16 cartridges and allows you to stack and expand the TL891/TL892 MiniLibraries.

^cThe TL893/896 Automated Tape Libraries also require Version 2B MUC firmware.

1.3 List of Subsumed Patches

You can find a list of patches for reported problems that have been corrected in this release on the Tru64 UNIX Operating System Associated Products Volume 2 CD-ROM, in the following file:

`<mountpoint>/TCR/doc/txt/TCR_V1.6_Subsumed_Patches.TXT`

1.4 KZPSA without Devices Causes Machine Check During Shutdown

If an AlphaServer 1200, 4000, or 4100 system in an Available Server or Production Server environment has a KZPSA SCSI adapter that is not attached to a device, a hardware problem causes a machine check during system shutdown.

Ensure that all KZPSA SCSI adapters on these systems are properly terminated and connected to a device. They do not have to be on a shared SCSI bus, just connected to a device that is powered on.

As an alternate measure, you can comment out the appropriate SCSI bus entry in the `/sys/conf/SYSTEM_NAME` system configuration file. For example:

```
⋮
bus          pza0          at pci3      slot 2 vector  pzaintr
controller   scsi8         at pza0      slot 0
bus          pza1          at pci3      slot 3 vector  pzaintr
controller   scsi9         at pza1      slot 0
bus          pza2          at pci3      slot 4 vector  pzaintr
controller   scsi10        at pza2      slot 0
bus          pza3          at pci3      slot 5 vector  pzaintr
controller   scsi11        at pza3      slot 0
bus          pci2          at mcbus0    slot 6
bus          pza4          at pci2      slot 2 vector  pzaintr
controller   scsi12        at pza4      slot 0
bus          pza5          at pci2      slot 3 vector  pzaintr
controller   scsi13        at pza5      slot 0
bus          pza6          at pci2      slot 4 vector  pzaintr
controller   scsi14        at pza6      slot 0
#bus         pza7          at pci2      slot 5 vector  pzaintr
#controller   scsi15        at pza7      slot 0
⋮
```

1.5 Do Not Use `ifconfig` or `nfs_ifconfig` to Configure Multiple Addresses for an NFS Service

To ensure correct NFS locking for NFS services with multiple client networks, do not use the `/var/ase/sbin/nfs_ifconfig` script, or the `ifconfig` command in the start scripts of an NFS service to configure multiple addresses.

A service is associated with one `rpc.lockd` daemon, which coordinates NFS locking for that service. If multiple addresses are configured for the service, and a request comes in for another IP address, it will get a lock from a different `rpc.lockd` daemon. Lock synchronization is not guaranteed between the daemons.

Instead, use the `ASEROUTING` configuration variable, which advertises host-based routes for an ASE service from a server that has multiple network interfaces. The use of the `ASEROUTING` configuration variable is described in Chapter 3 of the *TruCluster Software Products Administration* manual.

1.6 The use of `ifconfig down` is not a Replacement for Cable Pulls

If acceptance tests are used to ensure correct ASE behavior, the correct procedure is to pull the network cable. Do not use the `ifconfig interface_id down` command to simulate an interface going bad. The behavior may not be the same as pulling the network cable.

1.7 Tuning NFS Retransmission Timeout

If a UDP client of an ASE NFS service has not received a response to an NFS request, and has reached the maximum number of retries, the request will be retried according to the `mount` command `maxtimo` option. The default value is 20 seconds. If an ASE service relocates, the 20-second timeout may be too long for some clients to wait for the NFS request to be retried.

If this is the case, use the `mount -o maxtimo=n mount` option to decrease the time between request transmissions.

1.8 The `cpq_mibs` Daemon May Generate Binary Error Log Entries in an ASE Environment

When you install the Tru64 UNIX operating system subsets, you install Insight Manager, a Compaq application that allows a user with a Windows 95 or Windows NT personal computer to monitor and configure remote systems over the Internet or Intranet.

The `cpq_mibs` (`/usr/sbin/cpq_mibs`) and `insightd` (`/usr/share/sysman/bin/insightd`) daemons are used by the Insight Manager. The daemons are started as part of the system startup. The `insightd` daemon is started by the `/sbin/init.d/insightd` script. The `cpq_mibs` daemon is started by the `/sbin/init.d/snmpd` script along with the `snmpd` and `os_mibs` daemons.

The `cpq_mibs` daemon probes every device in the `/dev` directory when the system boots. For every disk reserved by a TruCluster Available Server Software or TruCluster Production Server Software highly available service, an error is written to the binary error log (CAM SCSI error) when the `cpq_mibs` daemon attempts to probe the device and finds it reserved.

If your TruCluster configuration consists of multiple NFS or disk services with a large number of disks, the binary error log could contain a large number of CAM SCSI errors.

If this is a concern to you, you can prevent these error log entries by preventing the `cpq_mibs` daemon from starting.

Note

There is no configuration file for the `cpq_mibs` daemon, so there is no way to tell it to ignore certain disks.

If you wish to prevent the `cpq_mibs` daemon from starting at system boot, edit the `/sbin/init.d/snmpd` script and comment out the lines used to start the `cpq_mibs` daemon as shown in Example 1-1.

Example 1-1: Preventing the `cpq_mibs` Daemon From Starting

```
:
:
    set `who -r`
    if [ $9 = "S" ] && [ "X$MasterPid" = "X" ]
then
    [ -f $SNMPD ] && {
    $SNMPD
        echo "Extensible SNMP master agent started"
    if [ "X$SubagentPid" = "X" ]
    then
        $OSMIBS &
    echo "Base O/S sub-agent started"
    fi
    if [ "X$SvrSystemPid" = "X" ]
    then
    $SVRSYSMIB > /dev/console 2>&1 &
        echo "Server System sub-agent started"
    fi
    if [ "X$SvrMgtPid" = "X" ]
    then
    $SVRMGMTMIB > /dev/console 2>&1 &
        echo "Server Management sub-agent started"
    fi
    # if [ "X$CpqMibsPid" = "X" ]
    # then
    # $CPQMIBS > /dev/console 2>&1 &
    #     echo "Compaq Management sub-agent started"
    # fi
    }
    fi
    ;;
:
:
```

1.9 Memory Channel Notes and Restrictions

The following notes apply to the cluster Memory Channel subsystem in TruCluster Production Server Software and TruCluster Memory Channel Software configurations only:

- One of the major differences between the MC1 and MC2 hardware is the use of Memory Channel address space. MC1 supports 128-MB Memory Channel address space and MC2 supports either 128-MB or

512-MB Memory Channel address space. You cannot mix MC1 and MC2 modules on the same rail, but you can mix MC1 and MC2 rails in the same cluster. Memory Channel address space is additive for multiple active rails, used for TruCluster Memory Channel Software, and restricted to the lesser of the two address spaces in a failover pair (used with TruCluster Production Server Software) with both MC1 and MC2 rails in the failover pair.

- If you install MC2 adapters in a cluster, the default is to use 128-MB address space even if the jumpers are set to select 512-MB address space. You have to enable the address space increase from 128-MB to 512-MB.

If you upgrade a cluster from MC1 to MC2 hardware, you cannot increase the address space until all MC1 adapters have been replaced with an MC2 module, and then you have to promote the address space increase from 128-MB to 512-MB.

Most of the Memory Channel restrictions are provided in Chapter 5 of the TruCluster Software Products *Hardware Configuration* manual. The following restrictions are provided:

- If you will be installing two Memory Channel 2 adapters in the DWLPB PCIA option of an AlphaServer 8x00, and promoting them to 512-MB PCI address space, do not install any other devices. The two Memory Channel 2 adapters consume the entire DWLPB PCI address space.
- The DS20 and ES40 only support Memory Channel 2 PCI adapters. They do not support Memory Channel 1 or Memory Channel 1.5 adapters.

1.10 Fibre Channel Requirements and Restrictions

Table 1–9 shows the supported AlphaServer systems with Fibre Channel and the number of KGPSA-BC PCI-to-Fibre Channel adapters supported on each system.

Table 1–9: AlphaServer Systems Supported for Fibre Channel

AlphaServer	Number of KGPSA-BC Adapters Supported
AS800	2
AS1000A	2
AS1200	4
AS4x00	4

Table 1–9: AlphaServer Systems Supported for Fibre Channel (cont.)

AlphaServer	Number of KGPSA-BC Adapters Supported
AS8x00 ^a	2
GS60 ^a	2
GS140 ^a	2

^aThe KGPSA-BC PCI-to-Fibre Channel adapter is only supported on the DWLPB PCIA option; it is not supported on the DWLPA.

The following requirements and restrictions apply to the use of Fibre Channel with the TruCluster Products:

- The root (/) and swap filesystems are not supported with Fibre Channel storage.
- A maximum of four member systems is supported
- The only supported Fibre Channel adapter is the KGPSA-BC PCI-to-optical Fibre Channel host adapter. It comes with an optical gigabit link module (GLM).
- The TruCluster Software products only support single-switch fabric configurations.
- The only Fibre Channel switch supported is the 8/16 Port DSGGA Fibre Channel switch.
- The DSGGA Fibre Channel switch supports both 50 and 62.5 micron multi-mode fibre (MMF) with the standard SC connector. The gigabit interface converter (GBIC) module supported is the short wavelength fibre optic version; therefore, only the 50 micron MMF optical cable is supported.
- Fibre Channel arbitrated loop (FC-AL) is not supported.
- The Fibre Channel RAID Array 8000 (RA8000) midrange departmental storage subsystem and Fibre Channel Enterprise Storage Array 12000 (ESA12000) house two HSG80 dual-channel controllers. There are provisions for 6 UltraSCSI channels.
- Only disk devices attached to the HSG80 Fibre Channel to Six Channel UltraSCSI Array controller are supported.
- No tape devices are supported.
- A maximum of eight targets is supported on each bus, with a switch representing a bus.
- A storage array with dual-redundant HSG80 controllers is two targets and consumes 4 ports on a switch.

- Each KGPSA is one target.
- A host is only allowed to have one KGPSA connected to each switch.
- The HSG80 only supports transparent failover mode when used in a TruCluster Software Products configuration.

1.11 UltraSCSI Hub Restrictions

The DS-DWZZH-03 and DS-DWZZH-05 UltraSCSI hubs are only supported on the following systems:

- AS800
- AS1000A
- AS1200
- AS4x00
- AS8x00
- GS60
- GS140

1.12 Enhanced Security in an Available Server or Production Server Configuration

The following discussion is based on an Available Server or Production Server configuration consisting of Compaq AlphaServer systems running Tru64 UNIX Version 4.0F (with enhanced security enabled) and TruCluster Available Server Software Version 1.6 or TruCluster Production Server Software Version 1.6.

1.12.1 Issues with Distributed Logins, Enhanced Security, and NIS

Tru64 UNIX with enhanced security enabled supports the NIS distribution of the extended user profiles maintained in the protected password database. This protected password database is used in conjunction with the `/etc/passwd` file to store user account information.

To minimize use of the `root` file system, the local extended user profiles are stored in two locations:

- User profiles for UIDs 0-99 are stored in `/tcb/files/auth.db`.
- User profiles for UIDs starting at 100 are stored in `/var/tcb/files/auth.db`.

These extended user profiles can be distributed by NIS as the `prpasswd` map, in the same manner that `/etc/passwd` can be distributed as the `passwd` map.

Note

See `prpasswd(4)` and `authcap(4)` for information on user profiles and the protected password (authentication) database.

Customers need secure, distributed, highly available logins. At first glance, the simplest solution is to enable enhanced security, install the TruCluster Available Server Software or TruCluster Production Server Software product, create an Available Server Environment (ASE), and then use NIS within the ASE to maintain and share a common authentication database.

However, in a Tru64 UNIX Version 4.0F enhanced security NIS environment, each user login requires an update of the `prpasswd` map on the NIS master and a successful push to all slave servers. Enhanced security requires that security-critical information in all copies of the authentication profile be identical across all servers; should the master fail, each slave must have a guaranteed correct copy of the `prpasswd` map. A side effect of these increased authentication requirements is a defacto single point of failure, the NIS master. Therefore, a highly available login service for enhanced security should not rely on NIS serving the `prpasswd` map.

1.12.2 An ASE Solution Using NIS and a Highly Available NFS Service

The recommended method for providing distributed, secure, highly available logins employs an ASE configuration using both NIS and NFS is as follows:

- NIS serves the `passwd` and `group` maps (and any other maps except `prpasswd`). One ASE member is configured as the NIS master and all other ASE members are configured as NIS slaves. The NIS highly available service setup can be found in the `/var/ase/examples_supported/README` file.
- An ASE NFS highly available service is created to make the `/var/tcb/files/auth.db` database available to all ASE members. The IP address for this service is on the cluster interconnect subnet, and the service is available only to ASE members.

This configuration gives each ASE member NFS access to the shared `/var/tcb/files/auth.db` database, which replaces a potential single point of failure with a highly available NFS service.

The ASE NFS service for the `/var/tcb/files` directory requires at least one shared disk, on a shared SCSI bus. Even though an `auth.db` database supporting 25,000 users should fit on a normal-sized "a" partition, the entire disk must be available to the service to allow for service failover.

This means that the entire disk should be used for highly available NFS served data. If such a service already exists, you can use that existing NFS service to distribute the `/var/tcb/files` directory. If no such service exists, use the "c" partition of the shared disk.

Note

If a "no single point of failure" configuration is required, the data is mirrored across two disks each on a different shared SCSI bus. Use LSM to create a mirror consisting of the two disk "c" partitions.

Because the NFS shared file system containing `/var/tcb/files/auth.db` file requires root access, the file system must be exported with the `-root=hostname[:hostname]...` option listing only the ASE members. Likewise access to the cluster interconnect subnet must be restricted. Add access filter entries to `/etc/ifaccess.conf` to restrict subnet access to ASE members.

Because the `/var/tcb/files/auth.db` database contains account information associated with UIDs starting with 100, UIDs 0-99 are local to each system. Each member's local `/etc/passwd` file must contain the `+` NIS syntax after UID 99. All user accounts starting with UID 100 should be created and maintained only on the NIS master. The `/etc/passwd` account information for accounts starting with UID 100 is moved to `/var/yp/src/passwd` and the `passwd` map created and pushed to the NIS slave servers.

Note

The `/etc/passwd` and `/etc/group` files are distributed by NIS because they do not require updates during a user login. However, these maps do require updates to add, modify, or delete a user. Consequently, if the NIS master is down, the add, modify, and delete user operations are disabled.

Review the following documentation before attempting to set up this configuration:

- *Tru64 UNIX Version 4.0F Installation Guide and Release Notes*
- *Tru64 UNIX System Administration (LSM and AdvFS)*
- *Tru64 UNIX Network Administration (NFS and NIS)*
- *Tru64 UNIX Security*

- TruCluster Software Products *Hardware Configuration, Software Installation, Administration, and Release Notes*
- NIS highly available service setup in the `/var/ase/examples_supported/README` file.

1.12.3 Configuring Distributed Enhanced Security with the TruCluster Software Products

Ensure that all PAKs (Product Authorization Keys) and software are installed and configure as described in the following sections.

Note

The following procedure was developed to allow for the high availability of logins to a cluster with enhanced security prior to the development of NIS master failover. As of Version 1.6, NIS master failover is available and is the preferred method of setting up enhanced security in a cluster. NIS highly available service setup can be found in the `/var/ase/examples_supported/README` file.

1.12.3.1 Setting Up the Initial (NIS Master) System

To set up an NIS master, follow these steps:

1. Install Tru64 UNIX Version 4.0F on one member system.
2. Build the kernel with networking, BIND, LSM, NFS, AdvFS, enhanced security and any other required subsystems.
3. Configure the operating system using either `setup` or the individual setup scripts as follows:
 - a. Perform normal network configuration (`/usr/sbin/netsetup`) and then do the following:
 - i. Allocate the IP names and addresses for each member's cluster interconnect. The NFS service called `tcbauth` also needs an IP address on this cluster interconnect subnet.
 - ii. Put these IP names and addresses in the `/etc/hosts` file. To restrict access to the cluster interconnect subnet to ASE members, add entries to the `/etc/ifaccess.conf` file.

- iii. Ensure that all ASE members use the same routing daemon; for example, `routed`.
- b. Configure a name service (`/usr/sbin/bindsetup`).
- c. Configure NFS (`/usr/sbin/nfssetup`) to establish the system as a server. When prompted about locking, answer `yes`. When prompted about exporting directories, answer `yes` but do not specify any particular exports. (This starts the appropriate daemons.) Select other options as needed for your environment.
- d. Configure NTP (`/usr/sbin/ntpsetup`).
- e. Configure any other subsystems needed for your environment.
- f. Set up the Logical Storage Manager as follows:
 - i. Run the `volinstall` command to create special devices.
 - ii. Set up the root disk group `rootdg` as follows:
 - 1. Use the `disklabel` command to initialize the disk labels for disks associated with the `rootdg` whose disk labels have not been initialized. For example:


```
# /usr/sbin/disklabel -rw /dev/rrz5c rz28
```
 - 2. Use the `volsetup` command to initialize the root disk group, `rootdg`, on a disk that is local to the system, not shared (see the LSM documentation for more details). For example:


```
# /usr/sbin/volsetup rz5
```
- 4. Using the procedures in the TruCluster Software Products *Software Installation* manual, install TruCluster Available Server Software Version 1.6 or TruCluster Production Server Software Version 1.6.

Note

Do not run `asemgr` to populate the ASE member list at this time.

- Run the `/usr/sbin/clu_ivp -v` command to verify the installation.
- 5. Configure enhanced security (`/usr/sbin/secsetup`). Choose the default options. Enabling enhanced security requires a reboot which you should do at this time.
 - 6. Run the `/usr/sbin/nissetup` command and configure the system as the NIS master serving `passwd`, `group`, and any other maps

EXCEPT `prpasswd`. Select the most secure options for NIS. The `nissetup` utility does not configure `prpasswd` by default.

7. Configure the storage that will contain the shared `/var/tcb/files` directory.

The following instructions create an LSM mirror set to hold this directory. This example uses disks `rz10c` and `rz20c`, both RZ28s. The actual disk partition names and disk types may be different on your system. Also, this example assumes that `rz5c` is used on each ASE member to establish the LSM `rootdg` root disk group.

- a. Initialize the disk labels for disks associated with the new disk group for the NFS service that have not been initialized before. For example:

```
# /usr/sbin/disklabel -rw /dev/rrz10c rz28
# /usr/sbin/disklabel -rw /dev/rrz20c rz28
```

- b. Run `voldiskadd` to create the new shared disk group `authdg` and to add the two disks:

```
# /usr/sbin/voldiskadd rz10
```

When prompted, enter the disk group (`authdg`), verify the new disk group (`y`), and enter the disk name (`authdg01`). For example:

```
# /usr/sbin/voldiskadd rz20
```

When prompted, enter the disk group (`authdg`) and the disk name (`authdg02`).

- c. Create a single disk volume with `authdg01` and mirror it onto `authdg02` using the `volassist` command. Create volume `v1_auth` in disk group `authdg` with a size of 2 GB (2000m) on `authdg01`:

```
# /usr/sbin/volassist -g authdg make v1_auth 2000m \
  usetype=fsgen authdg01
```

Set up the mirror of volume `v1_auth` with `authdg02` (this takes about 15 minutes for a pair of RZ28s):

```
# /usr/sbin/volassist -g authdg mirror v1_auth authdg02
```

- d. Create an AdvFS domain, `authdom`, containing fileset `authfs` on the new mirrored LSM volume `v1_auth`:

```
# /usr/sbin/mkfdmn /dev/vol/authdg/v1_auth authdom
# /usr/sbin/mkfset authdom authfs
```

8. Set up the ASE and the NFS service by running the `/usr/sbin/asemgr` program:

- a. Add the `tcbauth` NFS service:

```

service name = tcbauth
Advfs fileset = authdom#authfs
exported path = /auth
export list = <only cluster member names>
access mode = read-write
mount options = <no special mount options>
service policy = balanced
relocate on boot= no

```

- b. Modify the exports for the `tcbauth` service. Add the following exports entry option using your own ASE member cluster interconnect names:

```
/auth node0 node1 node2 -root=node0:node1:node2
```

9. Set up a highly available NIS service as described in the `/var/ase/examples_supported/README` file (preferred method) or set up a user-defined service on each member that will run the machine-specific start action script. This is required to allow a service to determine that the favored machine has failed.

- a. Sample ASE manager inputs are as follows:

```

service name = storm_hes
Advfs fileset =
exported path =
export list =
access mode =
mount options =
service policy = favored <hostname>
relocate on boot = yes

```

- b. Make a copy of the start action script example for each member giving each script a unique name, for example `storm_hes`.
 - c. Edit the start action scripts to replace `storm` with the node name of the machine that the start action script is to run on.
 - d. Place a copy of the compiled `hes` command on each member in the `/usr/sbin/` directory.
10. Run `nfssetup`. At the Add any remote file systems to be mounted prompt, mount the NFS served file system from `tcbauth` on `/auth`.
 11. Check that `/auth` is indeed mounted and check your ASE configuration with the following command:

```
# asemgr -d
```

1.12.3.2 Setting Up Additional Member Systems

Choose the next member to be installed and perform the following steps:

1. Install Tru64 UNIX Version 4.0F on the system.
2. Build the kernel with networking, BIND, LSM, NFS, AdvFS, enhanced security, and any other required subsystems.
3. Configure the system using either `setup` or the individual setup scripts as follows:
 - a. Perform normal network configuration (`/usr/sbin/netsetup`) and then do the following:
 - i. Allocate the IP names and addresses for each member's cluster interconnect subnet. The NFS service called `tcbauth` also needs an IP address on this cluster interconnect subnet.
 - ii. Put these IP names and addresses in the `/etc/hosts` file. To restrict access to the cluster interconnect subnet to ASE members, add entries to the `/etc/ifaccess.conf` file.
 - iii. Ensure that all ASE members use the same routing daemon; for example, `routed`.
 - b. Configure a name service (`/usr/sbin/bindsetup`).
 - c. Run `/usr/sbin/nfssetup` to configure NFS and establish the system as a server. When prompted about locking, answer `yes`. When prompted about exporting directories, answer `yes` but do not specify any particular exports. (This starts the appropriate daemons.) Select other options as needed for your environment.
 - d. Configure NTP (`/usr/sbin/ntpsetup`).
 - e. Configure any other subsystems needed for your environment.
 - f. Set up the Logical Storage Manager as follows:
 - i. Run the `volinstall` command to create special devices.
 - ii. Set up the root disk group `rootdg` as follows:
 1. Use the `disklabel` command to initialize the disk labels for disks associated with the `rootdg` whose disk labels have not been initialized. For example:

```
# /usr/sbin/disklabel -rw /dev/rrz5c rz28
```
 2. Use `volsetup` to initialize the root disk group, `rootdg`, on a disk that is local to the system, not shared (see the LSM documentation for more details). For example:

```
# /usr/sbin/volsetup rz5
```

4. Using the procedures in the TruCluster Software Products *Software Installation* manual, install TruCluster Available Server Software Version 1.6 or TruCluster Production Server Software Version 1.6.

Note

Do not run `asemgr` to populate the ASE member list at this time.

Run the `/usr/sbin/clu_ivp -v` command to verify the installation.

5. Configure enhanced security (`/usr/sbin/secsetup`). Choose the default options. Reboot to Enable enhanced security. at this time.
6. Run `nissetup` to configure this system as a slave server in the same domain as the master which was previously configured.
7. Run `nfssetup`. At the Add any remote file systems to be mounted prompt, mount the NFS served file system from `tcbauth` on `/auth`.
8. Install and configure any other members by repeating this procedure.

After all members are installed, run the ASE manager (`asemgr`) from the first ASE member installed (configured in Section 1.12.3.1) and add the newly installed member into the ASE. See the TruCluster Software Products *Administration* manual for more information.

1.12.3.3 Completing the Distributed Accounts Set Up on the NIS Master

On the NIS master configured in Section 1.12.3.1, follow these steps:

1. Copy the current `auth.db` to the `/auth` NFS file system as follows:

```
# cd /var/tcb
# cp -p -R files /auth/files
```

2. Move and link the NFS database as follows:

```
# mv /var/tcb/files /var/tcb/files.orig
# ln -s /auth/files /var/tcb/files
```

3. To verify that the database is operational, add a user account and try a login, password change, and a logout.
4. Add the shared user accounts starting with a UID greater than 100). This can be done with the `useradd` utility or scripts for large numbers of users.

Note

The enhanced security database format changed between DIGITAL UNIX Version 3.2* and Version 4.0* from a set of text files to an actual database (`auth.db`), so any scripts created prior to Version 4.0 to add users in an enhanced security environment have to be modified from creating/editing the text files to performing a series of `/tcb/bin/edauth` commands.

Some extra steps are required when using `useradd` because the utility will either add a user to the local or to the NIS distributed databases. In this setup, the system sees the extended authentication profile database as local (because the NFS disk is mounted locally), and the `passwd` and `group` information as distributed. After the user account has been added locally, the `passwd` information must be moved from the `/etc/passwd` to `/var/yp/src/passwd`, and the `group` information moved from `/etc/group` to `/var/yp/src/group`.

The NIS maps are then rebuilt and pushed. These steps only have to be done once per set of users being added; they do not have to be done separately for each user. Use the following procedure to add users individually, and to verify that the distributed enhanced security and ASE configuration is operating properly:

1. Create a user account:

```
# useradd -c "test user1" -d /usr/users/test1 -g users \  
-u 779 -x "local=1" test1
```

The `useradd` command creates the following:

```
# /tcb/bin/edauth -d p -g test1
```

```
t1:u_name=t1:u_id#64000:u_exp#0:u_life#0: \  
:chkent:
```

2. Edit the extended profile for `test1` to unlock the account:

```
# /tcp/bin/edauth -d p test1
```

The `edauth` program invokes the `ed` editor. You need to edit the account and add the `u_lock` field followed by the `@ sign, u_lock@`, in order to unlock the account. For example, in `ed`:

```
# /tcb/bin/edauth -d p test1  
"/tmp/edaabyksa" 2 lines, 53 characters  
:/chk  
:chkent:  
:s/chkent/u_lock@:chkent/
```

```
:u_lock@:chkent:
:w
"/tmp/edaabyksa" 2 lines, 61 characters
:q
```

3. Add a password for the user:

```
# passwd test1
```

4. Using a text editor move the `test1` account information from `/etc/passwd/` to `/var/yp/src/passwd`.

5. Update the NIS maps:

```
# cd /var/yp
# make
```

6. As user `test1`, test the account by logging in, changing the password, and logging out.

1.12.3.4 Completing the Distributed Accounts Set Up on the Additional Member Systems

Complete the setup on additional members using the procedures in this section.

1. On the remaining cluster members, do the following:

```
# mv /var/tcb/files /var/tcb/files.orig
# ln -s /auth/files /var/tcb/files
```

2. Verify that all cluster members can change passwords, and login and logout the newly added user.

Note

You may have to change permissions on the `/auth` directory mount point to `775` to allow users added on the master to log in on the client.

1.12.4 Audit in an ASE Configuration

In a Version 1.6 ASE configuration, audit is run independently on each member of the ASE configuration.

1.12.5 Available Server and Production Server NFS File Lock Configuration

This section documents the NFS File Lock Configuration in a cluster environment with enhanced security running.

1.12.5.1 Purpose

If a member of an ASE with a mounted file system that is part of an NFS service crashes while holding a lock on a file in the file system, the lock is not released by the member system running the NFS service. The system can be configured to solve this issue. If you are distributing `auth.db` over NFS as described in this note, you must configure NFS file locks as described in the following section or enhanced security logins will hang on an NFS server failover.

1.12.5.2 Procedure

You can configure the system to have the ASE agent free the NFS locks in case of a crash. NFS locks can be freed for any one ASE NFS service or for all ASE NFS services by setting the

`ASE_FREENFSLOCKS_ase_service_name` or `ASE_FREENFSLOCKS /etc/rc.config` variable to `yes`.

For example, to free the NFS locks for NFS service `ase1`, set the `ASE_FREENFSLOCKS_ase1` variable on all systems that are members of the ASE as follows:

```
# rcmgr set ASE_FREENFSLOCKS_ase1 yes
```

To free the NFS locks for all ASE NFS services, set the `ASE_FREENFSLOCKS` variable on all systems that are members of the ASE as follows:

```
# rcmgr set ASE_FREENFSLOCKS yes
```

Note

Use these procedures only when enhanced security is running.

The `ASE_FREENFSLOCKS /etc/rc.config` variable is an overriding variable, and if you use it you may free locks you did not intend to free.

1.12.6 Administration Notes

- The system administrator should add, modify, and delete user accounts on the NIS master.
- Instruct users to use the `passwd` (not `yppasswd`) command to change their passwords.

Index

A

accessing device special files, 1-4t
action script timeout values, 1-4t
advfsd daemon, 1-4t
AlphaServer systems
 supported, 1-5
AltaVista, 1-4
Apache web server, 1-4
Asynchronous Transfer Mode
 (*See* *ATM*)
ATM, 1-10
 Lan Emulation (LANE), 1-10

B

Baan, 1-4
busy mount point, 1-4t

C

cnxmibd daemon, 1-4t
commands
 fuser, 1-4t
 ifconfig, 1-13
 mount, 1-14
 /var/ase/sbin/nfs_ifconfig, 1-13
Common cluster MIB, 1-4t
configuration variables
 ASE_ADVFS ignore disks,
 1-4t
 ASE_IGNORE_BCL, 1-4t
 ASE_SNMPD_IGNORE_DISKS,
 1-4t
cpq_mibs daemon, 1-14
 errors caused by, 1-15
 preventing from starting, 1-15

D

daemons
 advfsd, 1-4t
 cnxmibd, 1-4t
 cpq_mibs, 1-14
 insightd, 1-14
 os_mibs, 1-4t
 rpc.lockd, 1-13
defining ASE services, 1-4t
disks
 required firmware revision, 1-10
documentation
 HTML versions of, vi
DRD, 1-4t
DS-DWZZH-03, 1-4t
DS-DWZZH-05, 1-4t

E

enhanced security, 1-4t
example scripts, 1-4t

F

Fibre Channel
 required firmware revisions, 1-7
 restrictions, 1-16
firmware requirements, 1-5
fuser, 1-4t

H

hardware requirements, 1-4
HTML
 documentation in, vi

I

ifconfig, 1-13
INFORMIX, 1-4
insightd daemon, 1-14

K

KZMSA
 required firmware revision, 1-6
KZPBA-CB, 1-4t, 1-12t
 required firmware revision, 1-6
KZPSA, 1-12t
 required firmware revision, 1-6
KZTSA
 required firmware revision, 1-6

M

maxtime mount option, 1-14
member systems
 required firmware revision, 1-5
 supported, 1-5
 supported SCSI adapters, 1-5
Memory Channel, 1-4t
 address space, 1-4t
 supported hardware, 1-6
modify DRD device special file
 group settings, 1-4t
 owner settings, 1-4t
 permissions, 1-4t
monitored network interfaces, 1-4t
mount
 dual, 1-4t
mount option
 maxtime, 1-14

N

NetRAIN, 1-4t
Netscape, 1-4
network adapters, 1-4t, 1-9
 ATM, 1-10
 supported ATM adapters, 1-10

 supported Ethernet adapters,
 1-9
 new functionality, 1-1
NFS, 1-4t
 over TCP, 1-4t

O

Oracle, 1-4
os_mibs daemon, 1-4t

P

panicking ASE, 1-4t
PMAZC
 required firmware revision, 1-6

R

RAID controllers
 required firmware revisions, 1-8
rc.config
 (*See configuration variables*)
RDG, 1-4t
recovery
 panicking ASE, 1-4t
reliable datagram
 (*See RDG*)
required firmware
 Fibre Channel, 1-7
 KZMSA, 1-6
 KZPBA-CB, 1-6
 KZPSA, 1-6
 KZTSA, 1-6
 member systems, 1-5
 PMAZC, 1-6
 RAID controllers, 1-8

S

Samba PC fileserver, 1-4
scripts, supported
 NIS, 1-4t
 printer, 1-4t

- scripts, unsupported
 - AltaVista search, 1-4
 - Apache web server, 1-4
 - Baan, 1-4
 - INFORMIX, 1-4
 - Netscape, 1-4
 - Oracle, 1-4
 - Samba PC fileserver, 1-4
 - Tuxedo, 1-4
- SCSI disks
 - required firmware revision, 1-10
- SCSI signal converters
 - required hardware revisions, 1-9
- SCSI tapes
 - required firmware revision, 1-11
- SCSI-2 adapters
 - KZMSA, 1-6
 - KZPBA-CB, 1-6
 - KZPSA, 1-6
 - KZTSA, 1-6
 - PMAZC, 1-6
 - required firmware revisions, 1-6
- security, 1-4t
- service configuration files, 1-4t
- shared SCSI bus, 1-12t
- SNMP, 1-4t

T

- tape devices
 - required firmware revision, 1-11
 - TL890, 1-12t
 - TL891, 1-12t
 - TL892, 1-12t
 - TL893, 1-12t
 - TL894, 1-12t
 - TL895, 1-12t
 - TL896, 1-12t
 - TZ885, 1-12t
 - TZ887, 1-12t
- timeout values, 1-4t
- Tuxedo, 1-4

U

- UltraSCSI Hub
 - restrictions, 1-18

V

- V1.6 functionality, 1-4
- /var/ase/sbin/nfs_ifconfig, 1-13
- Version 1.6 new functionality, 1-1

How to Order Tru64 UNIX Documentation

You can order documentation for the Tru64 UNIX operating system and related products at the following Web site:

<http://www.businesslink.digital.com/>

If you need help deciding which documentation best meets your needs, see the Tru64 UNIX *Documentation Overview* or call **800-344-4825** in the United States and Canada. In Puerto Rico, call **787-781-0505**. 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/>

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-MT4AA-G8
Tru64 UNIX Documentation Kit	QA-MT4AA-GZ
End User Documentation Kit	QA-MT4AB-GZ
Startup Documentation Kit	QA-MT4AC-GZ
General User Documentation Kit	QA-MT4AD-GZ
System and Network Management Documentation Kit	QA-MT4AE-GZ
Developer's Documentation Kit	QA-MT5AA-GZ
General Programming Documentation Kit	QA-MT5AB-GZ
Windows Programming Documentation Kit	QA-MT5AC-GZ
Reference Pages Documentation Kit	QA-MT4AG-GZ
Device Driver Kit	QA-MT4AV-G8

Reader's Comments

TruCluster Software Products

Release Notes

AA-ROJAD-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
ZK03 3/Y32
110 SPIT BROOK RD
NASHUA NH 03062 9987



----- Do not cut or tear - fold here and tape -----

Cut on this line