TruCluster Software Products

Release Notes

Part Number: AA-R0JAC-TE

January 1998

Product Version:	TruCluster Production Server Software Version 1.5, TruCluster Available Server Software Version 1.5, and TruCluster MEMORY CHANNEL Software Version 1.5
Operating System and Version:	DIGITAL UNIX Version 4.0D

This manual provides important information about the TruCluster software products.

Digital Equipment Corporation
Maynard, Massachusetts

Contents

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 Production Server 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 DIGITAL 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

DIGITAL welcomes any comments and suggestions you have on this and other DIGITAL 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:

  Digital Equipment 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 DIGITAL UNIX that you are using.
- If known, the type of processor that is running the DIGITAL UNIX software.

The DIGITAL 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 DIGITAL technical support office. Information provided with the software media explains how to send problem reports to DIGITAL.

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.
file	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, `cat`(1) indicates that you can find information on the `cat` 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 1.5 of the TruCluster software products.

1.1 New Features

Version 1.5 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 Version 1.5

Feature	PS	MC	AS	For more information, see:
Support for AlphaServer 1200 systems.	Yes	Yes	Yes	*Hardware Configuration*
Support for the HSZ70 RAID controller.	Yes	N/A	Yes	*Hardware Configuration*
Support for tape devices on a shared SCSI bus with server failover for the POLYCENTER NetWorker Save and Restore product. This allows a system administrator to configure a single, highly available backup service for each ASE.	Yes	N/A	Yes	*Hardware Configuration, Administration*
Support for up to 8 member systems in a cluster (up to 4 per ASE).	Yes	Yes[a]	No[b]	*Hardware Configuration*
Support for failover across MEMORY CHANNEL interconnects in a redundant MEMORY CHANNEL configuration when their adapters are installed on different PCI buses on member systems.	Yes	Yes	N/A	*Hardware Configuration*

Table 1–1: New Features in Version 1.5 (cont.)

Feature	PS	MC	AS	For more information, see:
Support for multiple active logical rails for applications using the MEMORY CHANNEL application programming interfaces (API) library.	Yes[c]	Yes	N/A	Section 1.7 of these release notes, *MEMORY CHANNEL Application Programming Interfaces*
Other enhancements to the MEMORY CHANNEL API library.	Yes	Yes	N/A	Section 1.7 of these release notes, *MEMORY CHANNEL Application Programming Interfaces*
Support for asynchronous transfer mode (ATM) as a primary, backup, and client network in an Available Server configuration, and as a client network in a Production Server cluster	Yes	N/A	Yes	*Hardware Configuration*, *Administration*
Introduction of persistent resource locking in the distributed lock manager (DLM) to decrease recovery time for database applications.	Yes	N/A	N/A	`dlm_rd_attach`(3), `dlm_rd_collect`(3), `dlm_rd_detach`(3), `dlm_rd_validate`(3)
Extensions to the `asemgr` utility to display service and member information from the ASE database.	Yes	No	Yes	*Administration*, `asemgr`(8)
Ability to update a service's LSM or AdvFS storage configuration, as well as other service information (such as Automatic Service Placement policy and an NFS service's exports file), in the ASE database without interrupting a service's availability to its clients.	Yes	N/A	Yes	*Administration*

Table 1–1: New Features in Version 1.5 (cont.)

Feature	PS	MC	AS	For more information, see:
Improved performance of DRD remote read operations, as a result of peer-to-peer direct-memory access (DMA), in configurations in which the storage adapters and MEMORY CHANNEL hardware are on the same PCI segment.	Yes	N/A	N/A	*Administration*, `drd`(7), `drd_dma`(8)
Support for non-ASE member systems in a Production Server cluster.	Yes	N/A	N/A	*Software Installation*, *Administration*
Extensions to the Cluster Monitor	Yes	N/A	Yes	*Administration*, Cluster Monitor online help
Year 2000 readiness	Yes	Yes	Yes	The *DIGITAL UNIX Year 2000 Readiness* white paper (in HTML format) on the *DIGITAL UNIX V4.0D Documentation, Volume 1* CD-ROM

[a]Eight-member MEMORY CHANNEL API clusters have been supported in previous releases of the TruCluster MEMORY CHANNEL Software product. MEMORY CHANNEL API clusters do not contain ASEs
[b]The 4-member restriction still applies to the single ASE within an Available Server configuration.
[c]Production Server cluster operation supports only the failover pair mode of operation. Only applications using the MEMORY CHANNEL API library can utilize multiple active logical rails. This mode of operation does not allow for failover between logical rails.

You cannot take advantage of Version 1.5 features until the entire cluster has been upgraded to Version 1.5. For example, you cannot add a fifth member or add a non-ASE member to a cluster until the upgrade is complete. Please read the TruCluster Software Products *Software Installation* manual before installing any of the TruCluster software products.

Depending upon which TruCluster software product you are using, you should read certain sections of this manual, as follows:

- Users of the TruCluster Production Server Software should read the entire document.
- Users of the TruCluster MEMORY CHANNEL Software should read Section 1.1, Section 1.2, Section 1.6, and Section 1.7.
- Users of the TruCluster Available Server Software should read Section 1.1, Section 1.2, Section 1.4, Section 1.5, Section 1.8, Section 1.9, and Section 1.10.

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 1.5 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 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 1.5 software products, indicating which 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 DIGITAL UNIX Version 4.0D.

Table 1–2 also lists the minimum Alpha System Reference Manual (SRM) firmware version required for each system. For systems running DIGITAL UNIX Version 4.0D, firmware is distributed on the Alpha Systems Firmware Update Version 5.0 CD-ROM.

Table 1–2: Supported Member Systems

System	Minimum SRM Firmware for Version 4.0D	PS	MC	AS
DEC 3000	V7.0	No	No	Yes
DEC 7000 and DEC 10000	V1.7	No	No	Yes
AlphaServer 300	V6.6	No	No	Yes
AlphaServer 400	V6.6	No	No	Yes
AlphaServer 800	V5.0	Yes	Yes	Yes
AlphaServer 1000	V5.0	No	No	Yes
AlphaServer 1000A	V5.0	Yes	Yes	Yes
AlphaServer 1200	V5.0	Yes	Yes	Yes
AlphaServer 2000	V5.0	Yes	Yes	Yes
AlphaServer 2100	V5.0	Yes	Yes	Yes
AlphaServer 2100A	V5.0	Yes	Yes	Yes

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

System	Minimum SRM Firmware for Version 4.0D	PS	MC	AS
AlphaServer 4000	V5.0	Yes	Yes	Yes
AlphaServer 4000A	V5.0	Yes	Yes	Yes
AlphaServer 4100	V5.0	Yes	Yes	Yes
AlphaServer 8200 and 8400	V5.0	Yes	Yes	Yes

1.2.2 Supported MEMORY CHANNEL Hardware

Table 1–3 describes the MEMORY CHANNEL hardware supported by TruCluster Production Server Software Version 1.5 and TruCluster MEMORY CHANNEL Software Version 1.5. MEMORY CHANNEL hardware is not supported by TruCluster Available Server Software.

Table 1–3: Supported MEMORY CHANNEL Hardware

Hardware	Description	MEMORY CHANNEL Model Number
MEMORY CHANNEL adapter	Adapter used for a cluster interconnect connection.	CCMAA-AA or CCMAA-BA
MEMORY CHANNEL hub	PC-class enclosure that is populated with line cards and used to connect MEMORY CHANNEL adapters. A hub is required if you have more than two member systems.	CCMHA-AA
Line card	Line card for a MEMORY CHANNEL hub.	CCMLA-AA
Copper link cable	Cabling for adapter-to-hub and adapter-to-adapter (virtual hub) connections.	BC12N-10 (10 meters)

1.2.3 Supported SCSI-2 Adapters

Table 1–4 indicates which SCSI-2 adapters are supported by TruCluster Production Server Software Version 1.5 (PS) and TruCluster Available Server Software Version 1.5 (AS).

Table 1–4: Supported SCSI-2 Adapters

SCSI-2 Adapter	Hardware Revision	Firmware Revision	PS	AS
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

[a]You 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.

1.2.4 Supported RAID Controllers

Table 1–5 lists the RAID controllers supported by TruCluster Production Server Software Version 1.5 (PS) and TruCluster Available Server Software Version 1.5 (AS). It indicates which Hierarchical Storage Operating Firmware (HSOF) revisions are required for each RAID controller.

Table 1–5: Supported RAID Controllers

RAID Controller	Firmware Revision	PS	AS
HSZ10-A*x*	0306	No	Yes[a]
SWXRA-Z1 (HSZ20)	HSOF V3.0	Yes	Yes
HSZ40-A*x*	HSOF V2.0 or V2.5	Yes	Yes
HSZ40-B*x*	HSOF V2.5 or higher	Yes	Yes
HSZ40-C*x*	HSOF V2.7 or higher	Yes	Yes
HSZ50-A*x*	HSOF V5.0 or higher	Yes	Yes
HSZ70[b]	HSOF V7.0	Yes	Yes

[a]The HSZ10 controller can be used only in conjunction with PMAZC adapters in an Available Server configuration.
[b]The HSZ70 controller can be used only in conjunction with KZPSA adapters in a Production Server cluster or an Available Server configuration.

1.2.5 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

1.2.6 Supported Disks

Table 1–6 lists the disk devices supported by TruCluster Production Server Software Version 1.5 and TruCluster Available Server Software Version 1.5 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 1–6: 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
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

1.2.7 Supported Tapes

Table 1–7 lists the tape devices supported by TruCluster Production Server Software Version 1.5 and TruCluster Available Server Software Version 1.5 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.

Table 1–7: Supported Tape Devices

Tape Device	Firmware Revisions
TZ885	V52[a]
TZ88	V52
TZ89	V31

[a]Tape loader 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 DIGITAL UNIX Operating System Associated Products Volume 2 CD-ROM, in the following file:

<mountpoint>`//TCR/doc/txt/TCR_V1.5_Subsumed_Patches.TXT`

1.4 Network Adapters

The following notes apply to the use of network adapters in TruCluster Production Server clusters or TruCluster Available Server configurations:

- TruCluster Available Server Software and TruCluster Production Server Software do not support the operation of the DE500-XA and DE500-AA PCI/Ethernet network adapters in full-duplex mode, either at 10 Mbits/sec or 100 Mbits/sec. Both adapters are supported in half-duplex mode.
- Occasional system failures (in the form of `pcierror` panics) have been seen when local area transport (LAT) functionality is enabled on AlphaServer 8200 and 8400 systems that have DE435 PCI/Ethernet network adapters that are not connected to a network. You can use any of the following methods to resolve this problem:
  - Disable LAT functionality on the system by using the `latsetup` command.
  - Disable LAT functionality at the DE435 PCI/Ethernet network adapter by using the `latcp` command with the `–E` flag.
  - Attach each DE435 PCI/Ethernet network adapter on the system to a network.

1.5 Network File System Services

The following note applies to the use of Network File System (NFS) services in an ASE:

- Version 1.5 of TruCluster Production Server Software and TruCluster Available Server Software do not support NFS mounts over TCP. If a client has established an NFS mount to an ASE service over TCP (for example, `mount -t nfs -proto tcp`), a server failover may cause the client to hang.

  You can use the `nfsconfig` utility on each cluster member system to set the number of server TCP daemons to zero, which forces the NFS service not to accept any NFS connection over TCP.

1.6 MEMORY CHANNEL

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

- In prior versions of the TruCluster Production Server Software and the TruCluster MEMORY CHANNEL Software, MEMORY CHANNEL controller numbers were determined solely by probe order (that is, the first controller probed would become `mchan0`, the second `mchan1`, and so on). Starting in Version 1.5 of these products, MEMORY CHANNEL controller numbers are derived from the number listed in the system configuration file (in `/usr/sys/conf/`*`NAME`*). This provides users with greater flexibility in configuring MEMORY CHANNEL hardware.

  If you need the older behavior, edit the system configuration file so that MEMORY CHANNEL controller entries are positioned in probe order, and rebuild the kernel. For example:

  ```
  controller mchan1 at pci0 slot 6
  controller mchan0 at pci0 slot 7
  ```

- Avoid disconnecting a MEMORY CHANNEL cable from a member system in a running cluster. Disconnecting the cable breaks the cluster interconnect, separating the system from the rest of the cluster. However, if a MEMORY CHANNEL cable is accidentally disconnected from a member system that uses redundant MEMORY CHANNEL interconnects, each member system (including the one from which the cable was disconnected) attempts to fail over to the secondary MEMORY CHANNEL interconnect to reestablish connection with the lost member system.

  The failover process is initiated only if all member systems have a secondary MEMORY CHANNEL interface configured and available. When the member systems complete the failover, reconnect the MEMORY CHANNEL cable to the MEMORY CHANNEL adapter of the system from which it was disconnected.

- A member system could display a panic message, similar to the following, if it experiences a MEMORY CHANNEL receive/transmit error that causes the single active MEMORY CHANNEL interconnect to fail:

```
rmerror_int: Error_count = 2 unit = 0
                        Err_reg = 0xffffffffa0000001 Node = 1
panic (cpu 0): rmerror_int: fatal error and no alternate mc to failover
```

When you see this panic message, ensure that the MEMORY CHANNEL cabling is fully connected and reboot the member system.

- When a cluster node deliberately crashes another node, the crashing node displays a machine check message that includes output similar to the following:

```
Machine Check SYSTEM Fatal Abort
   .
   .
   .
MC SHUTDOWN
MEMORY CHANNEL SoftC for board #0: interrupt count = 242
        240 notifications
        2 state changes
        0 dropped acks
        0 error Interrupts (13 processed)
NO ERRORS LOGGED
   .
   .
   .
```

A full description of the machine check is displayed only if the `mchan_debug` attribute is set to 1 in the `mchan` subsystem section of the `/etc/sysconfigtab` file.

1.7 MEMORY CHANNEL Application Programming Interface Library

The following new features of the MEMORY CHANNEL application programming interface library (API) are available in Version 1.5 of the TruCluster Production Server Software and TruCluster MEMORY CHANNEL Software products:

- Multiple active rails

  The MEMORY CHANNEL API now supports multiple active MEMORY CHANNEL logical rails. A user may now allocate, read from, or write to MEMORY CHANNEL regions on up to four logical rails. Two new API functions have been introduced to support this functionality: `imc_ckerrcnt_mr`(3) and `imc_rderrcnt_mr`(3). Use these functions instead of `imc_ckerrcnt`(3) and `imc_rderrcnt`(3) in new or recompiled code. See the TruCluster Production Server Software *MEMORY CHANNEL Application Programming Interfaces* manual and the appropriate reference pages for a detailed discussion of this functionality.

- Point-to-point attach

Unless a different transmission mode is explicitly specified, all MEMORY CHANNEL transactions are broadcast. Starting in the Version 1.5 product, you can attach for transmit to a particular MEMORY CHANNEL region in point-to-point mode. As a result, all writes to the region are to the specified host only. This should generate increased bandwidth in future MEMORY CHANNEL implementations. Use this in situations in which processes on only one host must read from a particular region. See the TruCluster Production Server Software *MEMORY CHANNEL Application Programming Interfaces* manual for additional information on this feature.

- Cluster change notification

  The new `imc_wait_cluster_event`(3) function allows a user to watch for certain events in the MC API cluster, such as nodes leaving or joining the cluster or new logical rails coming on line. See the TruCluster Production Server Software *MEMORY CHANNEL Application Programming Interfaces* manual and the `imc_wait_cluster_event`(3) reference page for further information on this feature.

- Optimization for specifying addresses to which MEMORY CHANNEL receive regions are to be attached

  The `IMC_USE_ADDR` flag to the `imc_asattach`(3) function allows you to specify the address at which to attach MEMORY CHANNEL receive regions when that address represents a hole in the current process's address space. See the `imc_asattach`(3) reference page for more information.

- Faster spinlocks

  The `imc_lkacquire`(3) and `imc_lkrelease`(3) functions have been improved to provide a faster lock implementation.

- New get cluster information function

  The new `imc_getclusterinfo`(3) function allows a user to retrieve certain information about a MEMORY CHANNEL API cluster: the number of members and their names, the number of logical rails, and a bitmask of active logical rails. This function supersedes the `imc_gethosts`(3) function. All new code should use the newer function instead of `imc_gethosts`(3). See the TruCluster Production Server Software *MEMORY CHANNEL Application Programming Interfaces* manual and the `imc_getclusterinfo`(3) reference page for additional details.

- Efficient copy function

  The new `imc_bcopy`(3) function allows the user to copy efficiently to MEMORY CHANNEL transmit regions. See the `imc_bcopy`(3) reference page for additional information.

- API initialization function

  All programs should invoke the new imc_api_init(3) function before calling any other MEMORY CHANNEL API library functions. Existing programs that do not call this function will continue to work. See the TruCluster Production Server Software *MEMORY CHANNEL Application Programming Interfaces* manual and the imc_api_init(3) reference page for more information.

- Monitoring the MEMORY CHANNEL API with the imcs(1) command

  The new imcs command lets the user to see details about MEMORY CHANNEL regions allocated through the MEMORY CHANNEL API, such as their keys, size, and owner. The command also displays information about MEMORY CHANNEL spinlocks. See the imcs(1) reference page for details.

- Online examples

  Every example published in the TruCluster Production Server Software *MEMORY CHANNEL Application Programming Interfaces* manual is supplied on line in the kit at /usr/examples/cluster/mc_*. One additional example, of the use of the imc_bcopy function, is provided at this location.

- IMC_VERSION constant

  Starting in TruCluster MEMORY CHANNEL Software Version 1.5 and TruCluster Production Server Software Version 1.5, the imc.h file defines the IMC_VERSION symbol to the current version of the MEMORY CHANNEL API library. For Version 1.5, its value is 0x150.

1.8 Informational Message from dxlsm and dtadvfs Utilities

When you use the dxlsm or dtadvfs utility in a TruCluster Production Server cluster or a TruCluster Available Server configuration, the utility issues an informational message on startup and upon any configuration change. These informational messages remind you to run the asemgr utility to update the ASE configuration database to reflect changes in storage associated with ASE services. These informational messages are optional; you can suppress them by specifying the -N command-line flag when invoking the dxlsm or dtadvfs utility.

1.9 ASEROUTING Option Cannot Be Used with New gated (PS and AS)

You can use the ASEROUTING configuration variable only with the old gated (ogated). (ogated is the default selection in the netsetup script.) If you use the ASEROUTING configuration variable when the new gated is

running on ASE members, all service operations will fail and error messages are entered in the daemon.log file.

This restriction existed, but was undocumented, in previous releases of TruCluster Production Server Software and TruCluster Available Server Software. In previous releases, service operations would not fail and no errors were logged.

1.10 Kernel Attributes That Must Be Set Consistently on All Members

You must ensure that the settings of the /etc/sysconfigtab attributes listed in Table 1–8 are the same on the new system as on all member systems.

Table 1–8: Kernel Attributes That Must Be Identical on All Member Systems

Kernel Attribute	Subsystem	TruCluster Software Configuration	Description
dochecksum	mcnet	Production Server cluster only	Enables or disables TCP/IP checksums. By default, checksums are on (dochecksum=1). Turning checksums off results in a slight increase in performance of TCP/IP over MEMORY CHANNEL. See the TruCluster Software Products *Administration* manual for more information.
rx_mapping_enabled	mcnet	Production Server cluster only	Enables or disables copy avoidance. By default, copy avoidance is off (rx_mapping_enabled=0). Turning copy avoidance on results in a slight increase in performance of TCP/IP over MEMORY CHANNEL. See the TruCluster Software Products *Administration* manual for more information.

Table 1–8: Kernel Attributes That Must Be Identical on All Member Systems (cont.)

Kernel Attribute	Subsystem	TruCluster Software Configuration	Description
`rm_rail_style`	`rm`	Production Server cluster only	Configures the reliability style of the MEMORY CHANNEL interconects on a cluster member. See the TruCluster Software Products *Administration* manual and the TruCluster Production Server Software *MEMORY CHANNEL Application Programming Interfaces* manual for more information.
`enable_extended_uids`	`proc`	Production Server cluster or Available Server configuration	Enables or disables extended UIDs in the base operating system. See the DIGITAL UNIX *Release Notes* for more information.

Failure to match the current configuration can cause one or more systems to panic when you attempt to add a new system to the cluster. Thus, if you want to change any of these attributes in a running cluster, you must perform the following steps on each member system:

1. Edit the appropriate `/etc/sysconfigtab` attribute on each member. See Appendix B of the TruCluster Software Products *Administration* manual for instructions for doing so.
2. Halt each member system.
3. Reboot all member systems.

If you must change these attributes and do not intend to perform a rolling upgrade from Version 1.4A to Version 1.5, consider changing them during a simultaneous upgrade procedure. See the TruCluster Software Products *Software Installation* manual for a description of the upgrade procedures.

Index

N

P

R

S

T

How to Order Additional Documentation

Technical Support

If you need help deciding which documentation best meets your needs, call 800-DIGITAL (800-344-4825) before placing your electronic, telephone, or direct mail order.

Electronic Orders

To place an order at the Electronic Store, dial 800-234-1998 using a modem from anywhere in the USA, Canada, or Puerto Rico. If you need assistance using the Electronic Store, call 800-DIGITAL (800-344-4825).

Telephone and Direct Mail Orders

Your Location	Call	Contact
Continental USA, Alaska, or Hawaii	800-DIGITAL	Digital Equipment Corporation P.O. Box CS2008 Nashua, New Hampshire 03061
Puerto Rico	809-754-7575	Local Digital subsidiary
Canada	800-267-6215	Digital Equipment of Canada Attn: DECdirect Operations KAO2/2 P.O. Box 13000 100 Herzberg Road Kanata, Ontario, Canada K2K 2A6
International	—	Local Digital subsidiary or approved distributor
Internal (submit an Internal Software Order Form, EN-01740-07)	—	SSB Order Processing – NQO/V19 *or* U.S. Software Supply Business Digital Equipment Corporation 10 Cotton Road Nashua, NH 03063-1260

Reader's Comments

TruCluster Software Products
Release Notes
AA-R0JAC-TE

Digital 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)	☐	☐	☐	☐
Clarity (easy to understand)	☐	☐	☐	☐
Organization (structure of subject matter)	☐	☐	☐	☐
Figures (useful)	☐	☐	☐	☐
Examples (useful)	☐	☐	☐	☐
Index (ability to find topic)	☐	☐	☐	☐
Usability (ability to access information quickly)	☐	☐	☐	☐

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

digital™

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

DIGITAL EQUIPMENT CORPORATION
UEG PUBLICATIONS MANAGER
ZKO3−3/Y32
110 SPIT BROOK RD
NASHUA NH 03062−9987

Do Not Cut or Tear − Fold Here

Cut on
Dotted
Line