

Tru64 UNIX 5.1A and TruCluster Server 5.1A

Patch Summary and Release Notes for Patch Kit-0003

August 2002

This manual describes the release notes and contents of Patch Kit-0003. It provides special instructions for installing individual patches.

For information about installing or removing patches, baselining, and general patch management, see the *Patch Kit Installation Instructions*.

© 2002 Hewlett-Packard Company

Microsoft®, Windows®, and Windows NT® are trademarks of Microsoft Corporation in the U.S. and/or other countries. Intel® and Pentium® are trademarks of Intel Corporation in the U.S. and/or other countries. Motif®, OSF/1®, The Open Group™, and UNIX® are trademarks of The Open Group in the U.S. and/or other countries. All other product names mentioned herein may be trademarks or registered trademarks of their respective companies.

Confidential computer software. Valid license from Compaq Computer Corporation, a wholly owned subsidiary of Hewlett-Packard Company, required for possession, use, or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

None of Compaq, HP, or any of their subsidiaries shall be liable for technical or editorial errors or omissions contained herein. The information in this document is provided "as is" without warranty of any kind and is subject to change without notice. The warranties for Compaq products are set forth in the express limited warranty statements accompanying such products. Nothing herein should be construed as constituting an additional warranty.

Contents

About This Manual

1 Release Notes

1.1	Patch Process Resources	1-1
1.2	Required Storage Space	1-1
1.3	Inclusion of Base Level in tar File Name	1-2
1.4	Problem with lockmode=4 on AlphaServers with QLogic SCSI Disk Controllers	1-2
1.5	Updates for Rolling Upgrade Procedures	1-2
1.5.1	Unrecoverable Failure Procedure	1-2
1.5.2	During Rolling Upgrade, Do Not Add or Delete OSF, TCR, or IOSWW Subsets	1-2
1.5.3	depord Warnings and cat Errors	1-3
1.5.4	Undoing a Rolling Patch	1-3
1.5.5	Ignore Message About Missing ladebug.cat File During Rolling Upgrade	1-3
1.5.6	clu_upgrade undo of Install Stage Can Result in Incorrect File Permissions	1-4
1.5.7	Missing Entry Messages Can Be Ignored During Rolling Patch ..	1-4
1.5.8	Relocating AutoFS During a Rolling Upgrade on a Cluster	1-4
1.6	Updates to sys_check	1-5
1.6.1	TMPDIR Variable	1-5
1.6.2	sys_check Version 125 Web Kit	1-6
1.7	When Taking a Cluster Member to Single-User Mode, First Halt the Member	1-6
1.8	Additional Steps Required When Installing Patches Before Cluster Creation	1-7
1.9	Problems with clu_upgrade switch Stage	1-7
1.10	Support for SDLT160 Tape Device	1-7
1.11	Release Note for Tru64 UNIX Patch 156.00	1-8
1.12	Release Note for Tru64 UNIX Patches 226.00 and 228.00	1-10
1.13	Release Note for Tru64 UNIX Patch 252.00	1-11
1.14	Release Note for Tru64 UNIX Patch 807.00	1-11
1.15	Release Notes for Tru64 UNIX Patch 737.00	1-17
1.15.1	Updates to sh, csh, and ksh	1-17
1.15.2	sh noclobber Option and > , >> Constructs Added	1-17
1.15.3	ksh noclobber Behavior Clarified	1-17
1.15.4	csh noclobber Behavior Clarified	1-18
1.15.5	Updated mkdir System Call and Command	1-18
1.16	Release Note for Tru64 Patch 504.00	1-18
1.17	Release Note for Tru64 UNIX Patch 846.00	1-21
1.17.1	Enabling the /dev/poll Function	1-21
1.17.2	Removal of Version-switched patch	1-22
1.17.3	New ee Attribute	1-22
1.17.4	lag(7) Reference Page Update	1-22
1.17.5	lagconfig(8) Reference Page Update	1-25
1.17.6	wol(8) Reference Page Update	1-28

1.18	Release Note for TruCluster Patch 9.00	1-30
1.19	Release Note for TruCluster Patch 95.00	1-32
1.20	Release Note for TruCluster Patch 202.00	1-32
1.20.1	Enablers for EVM	1-32
1.20.2	Rolling Upgrade Version Switch	1-32
1.20.3	Restrictions Removed	1-33

2 Summary of Base Operating System Patches

3 Summary of TruCluster Software Patches

Tables

2-1	Updated Base Operating System Patches	2-1
2-2	Summary of Base Operating System Patches	2-3
3-1	Updated TruCluster Software Patches	3-1
3-2	Summary of TruCluster Patches	3-2

About This Manual

This manual contains information specific to Patch Kit-0003 for the Tru 64 UNIX 5.1A operating system and TruCluster Server Software 5.1A products. It provides a list of the patches contained in each kit and describes the information you need to know when installing specific patches.

For information about installing or removing patches, baselining, and general patch management, see the *Patch Kit Installation Instructions*.

Audience

This manual is for the person who installs and removes the patch kit and for anyone who manages patches after they are installed.

Organization

This manual is organized as follows:

- Chapter 1 Contains the release notes for this patch kit.
- Chapter 2 Summarizes the Tru64 UNIX operating system patches included in the kit.
- Chapter 3 Summarizes the TruCluster software patches included in the kit.

Related Documentation

In addition to this manual, you should be familiar with the concepts and mechanisms described in the following Tru64 UNIX and TruCluster documents:

- Tru64 UNIX and TruCluster *Patch Kit Installation Instructions*
- Tru64 UNIX *Patch Kit Installation Instructions*
- `dupatch(8)` reference page
- Tru64 UNIX *Installation Guide*
- TruCluster Server *Cluster Installation*
- TruCluster Server *Cluster Administration*
- Release-specific installation documentation

Reader's Comments

We welcome 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, ZK03-3/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
ZK03-3/Y32
110 Spit Brook Road
Nashua, NH 03062-9987

Please include the following information along with your comments:

- The full title of this document.
- 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 technical support office. Information provided with the software media explains how to send problem reports.

Release Notes

This chapter provides important information that you need in order to work with the Tru64 UNIX 5.1A and TruCluster 5.1A Patch Kit-0003.

1.1 Patch Process Resources

We provide Web sites to help you with the patching process:

- To obtain the latest patch kit for your operating system and cluster:
<http://ftp1.support.compaq.com/public/unix/>
- To view or print the latest version of the *Patch Kit Installation Instructions* or the *Patch Summary and Release Notes* for a specific patch kit:
<http://www.tru64unix.compaq.com/docs/patch/>
- To visit our main support page:
<http://www.compaq.com/support/index.shtml>
- To visit the Tru64 UNIX homepage:
<http://www.tru64unix.compaq.com/>

1.2 Required Storage Space

The following storage space is required to successfully install this patch kit:

Base Operating System

- Temporary Storage Space
A total of ~250 MB of storage space is required to untar this patch kit. We recommend that this kit not be placed in the `/`, `/usr`, or `/var` file systems because doing so may unduly constrain the available storage space for the patching activity.
- Permanent Storage Space
Up to ~71 MB of storage space in `/var/adm/patch/backup` may be required for archived original files if you choose to install and revert all patches. See the *Patch Kit Installation Instructions* for more information.
Up to 73 MB of storage space in `/var/adm/patch` may be required for original files if you choose to install and revert all patches. See the *Patch Kit Installation Instructions* for more information.
Up to ~1346 KB of storage space is required in `/var/adm/patch/doc` for patch abstract and README documentation.
A total of ~176 KB of storage space is needed in `/usr/sbin/dupatch` for the patch management utility.

TruCluster Server

- Temporary Storage Space
A total of ~250 MB of storage space is required to untar this patch kit. We recommend that this kit not be placed in the `/`, `/usr`, or `/var` file systems

because doing so may unduly constrain the available storage space for the patching activity.

- **Permanent Storage Space**

Up to ~24 MB of storage space in `/var/adm/patch/backup` may be required for archived original files if you choose to install and revert all patches. See the *Patch Kit Installation Instructions* for more information.

Up to ~25 MB of storage space in `/var/adm/patch` may be required for original files if you choose to install and revert all patches. See the *Patch Kit Installation Instructions* for more information.

Up to ~1221 KB of storage space is required in `/var/adm/patch/doc` for patch abstract and README documentation.

A total of ~184 KB of storage space is needed in `/usr/sbin/dupatch` for the patch management utility.

1.3 Inclusion of Base Level in tar File Name

With this release, the name of the `tar` file containing the patch distribution has been expanded to include the baselevel for which this kit was built. This formerly internal baselevel number has become a common way of identifying kits. For complete information, see Section 1.3 of the *Patch Kit Installation Instructions*.

1.4 Problem with `lockmode=4` on AlphaServers with QLogic SCSI Disk Controllers

When an AlphaServer GS Series system with a QLogic SCSI disk controller is set to `lockmode=4`, the system may panic on boot. We will provide a fix for this in the near future.

1.5 Updates for Rolling Upgrade Procedures

The following sections provide information on rolling upgrade procedures.

1.5.1 Unrecoverable Failure Procedure

The procedure to follow if you encounter unrecoverable failures while running `dupatch` during a rolling upgrade has changed. The new procedure calls for you to run the `clu_upgrade -undo install` command and then set the system baseline. The procedure is explained in the *Patch Kit Installation Instructions* as notes in Section 5.3 and Section 5.6.

1.5.2 During Rolling Upgrade, Do Not Add or Delete OSF, TCR, or IOSWW Subsets

During a rolling upgrade, do not use the `/usr/sbin/setld` command to add or delete any of the following subsets:

- Base Operating System subsets (those with the prefix `OSF`).
- TruCluster Server subsets (those with the prefix `TCR`).
- Worldwide Language Support (WLS) subsets (those with the prefix `IOSWW`).

Adding or deleting these subsets during a roll creates inconsistencies in the tagged files.

1.5.3 depord Warnings and cat Errors

This release note explains `depord` warnings and `cat` errors displayed during a rolling upgrade with patches.

These warnings are only encountered if a rolling upgrade has been performed on the lead member, followed by the installation of patches on the lead member. When the remaining members perform the roll operation using the `clu_upgrade` roll command, a number of warning and error messages are displayed. The warning messages are from the `depord` command and state that the `.ctrl` file for patch subsets cannot be found.

These `depord` warnings are followed by error messages from the `cat` command stating that the `.inv` file for patch subsets cannot be opened. These warning and error messages are benign and can be ignored. The following is a sample of the warning and error messages that will be displayed:

```
depord: warning, no .ctrl file for "TCRPAT00008600520"
depord: warning, no .ctrl file for "TCRPAT00008400520"
depord: warning, no .ctrl file for "TCRPAT00008200520"
depord: warning, no .ctrl file for "TCRPAT00008000520"

... additional messages skipped ...

cat: cannot open
/var/cluster/members/{memb}/adm/update/tmpstaydir/instctrl/OSFPAT00000032520.inv
cat: cannot open
/var/cluster/members/{memb}/adm/update/tmpstaydir/instctrl/OSFPAT00000500520.inv

... additional messages skipped ...
```

1.5.4 Undoing a Rolling Patch

When you undo the stages of a rolling upgrade, the stages must be undone in the correct order. However, the `clu_upgrade` command incorrectly allows a user undoing the stages of a rolling patch to run the `clu_upgrade undo preinstall` command before running the `clu_upgrade undo install` command.

The problem is that in the install stage, `clu_upgrade` cannot tell from the `dupatch` flag files whether the roll is going forward or backward. This ambiguity allows a user who is undoing a rolling patch to run the `clu_upgrade undo preinstall` command without first having run the `clu_upgrade undo install` command.

To avoid this problem when undoing the stages of a rolling patch, make sure to follow the documented procedure and undo the stages in order.

1.5.5 Ignore Message About Missing `ladebug.cat` File During Rolling Upgrade

When installing the patch kit during a rolling upgrade, you may see the following error and warning messages. You can ignore these messages and continue with the rolling upgrade.

```
Creating tagged files.
.....
.....
*** Error ***
The tar commands used to create tagged files in the '/usr' file system have
reported the following errors and warnings:
    tar: lib/nls/msg/en_US.88591/ladebug.cat : No such file or directory
.....
*** Warning ***
The above errors were detected during the cluster upgrade. If you believe that
the errors are not critical to system operation, you can choose to continue.
If you are unsure, you should check the cluster upgrade log and refer
to clu_upgrade(8) before continuing with the upgrade.
```

1.5.6 `clu_upgrade undo` of Install Stage Can Result in Incorrect File Permissions

This note applies only when both of the following are true:

- You are using `installupdate`, `dupatch`, or `nhd_install` to perform a rolling upgrade.
- You need to undo the install stage; that is, to use the `clu_upgrade undo install` command.

In this situation, incorrect file permissions can be set for files on the lead member. This can result in the failure of `rsh`, `rlogin`, and other commands that assume user IDs or identities by means of `setuid`.

The `clu_upgrade undo install` command must be run from a nonlead member that has access to the lead member's boot disk. After the command completes, follow these steps:

1. Boot the lead member to single-user mode.
2. Run the following script:

```
#!/usr/bin/ksh -p
#
#   Script for restoring installed permissions
#
cd /
for i in /usr/.smbd./$(OSF|TCR|IOS|OSH)*.sts
do
    grep -q "_INSTALLED" $i 2>/dev/null && /usr/sbin/fverify -y <"${i%.sts}.inv"
done
```

3. Rerun `installupdate`, `dupatch`, or `nhd_install`, whichever is appropriate, and complete the rolling upgrade.

For information about rolling upgrades, see Chapter 7 of the *Cluster Installation* manual, `installupdate(8)`, and `clu_upgrade(8)`.

1.5.7 Missing Entry Messages Can Be Ignored During Rolling Patch

During the `setup` stage of a rolling patch, you might see a message like the following:

```
Creating tagged files.
.....
.....
.....

clubase: Entry not found in /cluster/admin/tmp/stanza.stdin.597530

clubase: Entry not found in /cluster/admin/tmp/stanza.stdin.597568
```

An `Entry not found` message will appear once for each member in the cluster. The number in the message corresponds to a PID.

You can safely ignore this `Entry not found` message.

1.5.8 Relocating AutoFS During a Rolling Upgrade on a Cluster

This note applies only to performing rolling upgrades on cluster systems that use AutoFS.

During a cluster rolling upgrade, each cluster member is singly halted and rebooted several times. The *Patch Kit Installation Instructions* direct you to manually relocate applications under the control of Cluster Application Availability (CAA) prior to halting a member on which CAA applications run.

Depending on the amount of NFS traffic, the manual relocation of AutoFS may sometimes fail. Failure is most likely to occur when NFS traffic is heavy. The following procedure avoids that problem.

At the start of the rolling upgrade procedure, use the `caa_stat` command to learn which member is running AutoFS. For example:

```
# caa_stat -t
Name           Type           Target      State      Host
-----
autofs         application    ONLINE     ONLINE     rye
cluster_lockd application    ONLINE     ONLINE     rye
clustercron   application    ONLINE     ONLINE     swiss
dhcp          application    ONLINE     ONLINE     swiss
named         application    ONLINE     ONLINE     rye
```

To minimize your effort in the procedure described as follows, it is desirable to perform the roll stage last on the member where AutoFS runs.

When it comes time to perform a manual relocation on a member where AutoFS is running, follow these steps:

1. Stop AutoFS by entering the following command on the member where AutoFS runs:

```
# /usr/sbin/caa_stop -f autofs
```

2. Perform the manual relocation of other applications running on that member:

```
# /usr/sbin/caa_relocate -s current_member -c target_member
```

After the member that had been running AutoFS has been halted as part of the rolling upgrade procedure, restart AutoFS on a member that is still up. (If this is the roll stage and the halted member is not the last member to be rolled, you can minimize your effort by restarting AutoFS on the member you plan to roll last.)

1. On a member that is up, enter the following command to restart AutoFS. (The member where AutoFS is to run, *target_member*, must be up and running in multi-user mode.)

```
# /usr/sbin/caa_startautofs -c target_member
```

2. Continue with the rolling upgrade procedure.

1.6 Updates to `sys_check`

This section describes updates to the `sys_check` command.

1.6.1 TMPDIR Variable

If the `TMPDIR` environment variable is not defined, then `sys_check -escalate` will always put the `escalate.tar` files in `/var/tmp` even if you specify an alternate directory. To work around this problem, you must first set and export the `TMPDIR` environment variable to the directory where you want `sys_check` to put the `escalate.tar` files. For example, if you want `sys_check` to put the `escalate.tar` files in `/var/adm`, then you must execute the following commands before running `sys_check -escalate`.

```
# ksh
# export TMPDIR=/var/adm
# sys_check -escalate
```

1.6.2 sys_check Version 125 Web Kit

The following information is for users who have installed `sys_check` Version 125 web kit or higher and are currently using that version of `sys_check` in the web kit as the system default version.

This patch kit contains `sys_check` Version 124. If you have already installed the `sys_check` Version 125 web kit or higher, then installing this patch kit will downgrade the version of `sys_check` that is being used by the system. However, you can easily set the system default back to the version of `sys_check` that you downloaded from the web by using the `/usr/sbin/use_sys_check` script. For example, type **`use_sys_check 125`** at the command line prompt to set `sys_check` Version 125 as the system default.

If you wish to delete the `sys_check` patch (that is, `sys_check` Version 124) then you should make sure that Version 124 is the system default version before deleting the patch. You can verify this by examining the output of the `sys_check -v` command. If 124.0 is not the default version, then you should run the `/usr/sbin/use_sys_check 124` command to set the system default version of `sys_check` to version 124. Setting the system default to 124 ensures that the Version 124 `sys_check` files get removed when the patch is deleted.

After you delete the patch, the system default version of `sys_check` will automatically be set to the version of `sys_check` that you downloaded from the web. This is because `dupatch` saves the symbolic links that point to the web kit location when the patch gets installed and will restore these symbolic links when the patch gets deleted.

If you delete the patch and the system default version is not set to 124, then Version 124 will remain on the system because `sys_check` Version 124 has been backed up by the web kit (for example, `/usr/sbin/sys_check.124.0`).

You will encounter problems if you delete the `sys_check` web kit and then delete this patch kit. This is because `dupatch` will restore the symbolic links to the web kit location when the patch is deleted. If you have deleted the web kit, then the symbolic links will point to non-existent files. You can fix this problem by re-installing the `sys_check` web kit.

1.7 When Taking a Cluster Member to Single-User Mode, First Halt the Member

To take a cluster member from multi-user mode to single-user mode, first halt the member and then boot it to single-user mode. For example:

```
# shutdown -h now
>>> boot -fl s
```

Halting and booting the system ensures that it provides the minimal set of services to the cluster and that the running cluster has a minimal reliance on the member running in single-user mode.

When the system reaches single-user mode, run the following commands:

```
# init s
# bcheckrc
# lmf reset
```

1.8 Additional Steps Required When Installing Patches Before Cluster Creation

This note applies only if you install a patch kit before creating a cluster; that is, if you do the following:

1. Install the Tru64 UNIX base kit.
2. Install the TruCluster Server kit.
3. Install the Version 5.1A Patch Kit-0003 before running the `clu_create` command.

In this situation, you must then perform three additional steps:

1. Run `versw`, the version switch command, to set the new version identifier:

```
# /usr/sbin/versw -setnew
```
2. Run `versw` to switch to the new version:

```
# /usr/sbin/versw -switch
```
3. Run the `clu_create` command to create your cluster:

```
# /usr/sbin/clu_create
```

1.9 Problems with `clu_upgrade` switch Stage

If the `clu_upgrade` switch stage does not complete successfully, you may see a message like the following:

```
versw: No switch due to inconsistent versions
```

The problem can be due to one or more members running `genvmunix`, a generic kernel.

Use the command `clu_get_info -full` and note each member's version number, as reported in the line beginning

```
Member base O/S version
```

If a member has a version number different from that of the other members, shut down the member and reboot it from `vmunix`, the custom kernel. If multiple members have the different version numbers, reboot them one at a time from `vmunix`.

1.10 Support for SDLT160 Tape Device

You must add the following entries in the `/etc/DDR.dbase` and then run `/sbin/DDR_config` for the new SDLT160 tape device to be recognized.

1. Add the following to `/etc/DDR.dbase`:

```
scsi_density_table_size = 0x4a

scsi_tape_density[0x42] = "density_code_42"    0      0
scsi_tape_density[0x43] = "density_code_43"    0      0
scsi_tape_density[0x44] = "density_code_44"    0      0
scsi_tape_density[0x45] = "density_code_45"    0      0
scsi_tape_density[0x46] = "density_code_46"    0      0
scsi_tape_density[0x47] = "density_code_47"    0      0
scsi_tape_density[0x48] = "131000_bpi"        131000 0
scsi_tape_density[0x49] = "190000_bpi"        190000 0

SCSIDEVICE
#
# Matches SDLT320
#
Type = tape
Name = "COMPAQ" "SDLT320"
#
#
```

```

PARAMETERS:
  TypeSubClass      = tk
  TagQueueDepth    = 0
  MaxTransferSize   = 0x0fffffb      # (16MB - 4)
  ReadyTimeSeconds = 120             # seconds

DENSITY:
#
  DensityNumber = 0
  DensityCode = 0x48
  CompressionCode = 0x1
  Buffered = 0x1

DENSITY:
#
  DensityNumber = 1,5
  DensityCode = default
  CompressionCode = 0x1
  Buffered = 0x1

DENSITY:
#
  DensityNumber = 2,4,6,7
  DensityCode = default
  CompressionCode = 0x0
  Buffered = 0x1

DENSITY:
#
  DensityNumber = 3
  DensityCode = 0x48
  CompressionCode = 0x0
  Buffered = 0x1

```

2. Run `/sbin/DDR_config` (see `DDR_config(8)` for more information).

1.11 Release Note for Tru64 UNIX Patch 156.00

This release note updates the `envconfig(8)` reference page.

`envconfig(8)`

NAME

`envconfig` - Configures the Environmental Monitoring daemon

SYNOPSIS

`/usr/sbin/envconfig -c var=value`

`/usr/sbin/envconfig start | stop`

`/usr/sbin/envconfig -q`

OPTIONS

Environmental Monitoring provides a means of detecting system threshold conditions, that if exceeded, could result in a loss of data or damage to the system itself. To detect and notify users of critical conditions, the `envmond` daemon is used. This utility, `envconfig`, is used to customize the `envmond` daemon. This section describes the `envconfig` options you can use to configure the daemon.

`-c var=value`

Sets the variables that specify how the system environment is monitored. These variables are stored in the `/etc/rc.config` file and are read by the `envmond` daemon at system startup. If a variable is not set, the default value of that variable is assumed.

ENVMON_CONFIGURED

Specifies the state of Environmental Monitoring. If this variable is set to zero (0), the Environmental Monitoring package is not started during the system boot. If this variable is set to 1, and Environmental Monitoring is supported by that platform, it is

started during the system boot. The default value is zero (0).

ENVMON_GRACE_PERIOD

Specifies the time (in minutes) that can elapse between the detection of a high temperature condition and the shutdown of the system. The default value is 15 minutes.

ENVMON_HIGH_THRESH

Specifies the threshold level that can be encountered before the envmond daemon broadcasts a warning and suggested action.

ENVMON_MONITOR_PERIOD

Specifies the frequency (in seconds) between queries of the system by the envmond daemon. The default value is 60 seconds.

ENVMON_USER_SCRIPT

Specifies the path of a user-defined script that you want the envmond daemon to execute when a high threshold level is encountered. The envmond daemon continues to check the environment after the script has executed and proceeds as needed should the high threshold levels persist.

If you set this variable, the envmond daemon directs output from the script to /dev/console. Output is not displayed on standard output or written to a file as this is not the behavior of the daemon. To display on standard output, explicitly specify the logger command within the user defined script

ENVMON_SHUTDOWN_SCRIPT

Specifies the path of a user-defined shutdown script that you want the envmond daemon to execute when a shutdown condition is encountered. The envmond daemon will execute this script in place of /sbin/shutdown. If you want the system to be shut down and you configure a script for ENVMON_SHUTDOWN_SCRIPT you must execute /sbin/shutdown from within your script. If you do not specify anything for ENVMON_SHUTDOWN_SCRIPT envmond will, by default, run /sbin/shutdown when a shutdown condition is encountered.

If you set this variable, the envmond daemon directs output from the script to /dev/console. Output is not displayed on standard output or written to a file as this is not the behavior of the daemon. To display on standard output, explicitly specify the logger command within the user-defined script.

start | stop

Turns the envmond daemon on or off after system startup.

-q Displays the values of ENVMON_CONFIGURED, ENVMON_GRACE_PERIOD, ENVMON_HIGH_THRESH, ENVMON_MONITOR_PERIOD, ENVMON_USER_SCRIPT, and ENVMON_SHUTDOWN_SCRIPT as specified in the /etc/rc.config file. If a specified entry is not found, the environmental variable is not displayed.

DESCRIPTION

The envconfig utility is used to customize the envmond daemon. You must have root privileges to use this utility. Using this utility, you can:

- + Specify whether or not Environmental Monitoring is turned on or off at system startup.
- + Specify how much time can elapse between the envmond daemon encountering a critical condition and the daemon initiating an orderly shutdown of the system.
- + Specify how frequently the envmond daemon queries the system for information.
- + Start and stop the envmond after Environmental Monitoring has been turned on at system startup.
- + Display the settings of the environment variables as specified in the

/etc/rc.config file.

Note that the feature that you want to monitor must be supported on a given platform. For example, the AlphaServer 8400/GS140 supports reporting of power supply and fan status, the current system temperature, and the maximum allowed system temperature.

EXAMPLES

The following procedure describes how you test for and start the environmental monitoring subsystem

1. In multiuser mode, check the status of the environmental monitoring subsystem as follows:

```
# /sbin/sysconfig -q envmon
envmon:
env_current_temp = 35
env_high_temp_thresh = 40
env_fan_status = 0
env_ps_status = 0
env_supported = 1
```

2. If the value of `env_supported` is 0, configure the `envmond` daemon and reboot the system using either of the following methods:

- + At the command prompt, enter the following command:
/usr/sbin/envconfig -c ENVMON_CONFIGURED=1
- + Use the `rcmgr` command as follows:
rcmgr set ENVMON_CONFIGURED 1

This command will enable the `envmond` daemon and export the variable, creating the following two lines in the `/etc/rc.config` file:

```
ENVMON_CONFIGURED="1"
export ENVMON_CONFIGURED
```

You can use the `/sbin/sysconfig` command to view the system environment at any time. The `envmond` daemon will print warning messages in the event of a power supply failure, abnormality, or high temperatures. Error logs are logged in the `/var/adm/binary.errlog`.

In the following example, the system shuts down in 10 minutes if the temperature does not fall below the critical threshold.

```
/usr/sbin/envconfig -c ENVMON_GRACE_PERIOD=10
```

FILES

`/etc/rc.config*`

Databases that contain the values of the environment monitoring variables. Note that you must use the `rcmgr` command to update the `rc.config*` files, particularly on clustered systems.

SEE ALSO

Commands: `envmond(8)`

1.12 Release Note for Tru64 UNIX Patches 226.00 and 228.00

Patches 226.00 and 228.00 deliver version V2.0-094d of the `libots3` library. If your system has the Compaq FORTRAN Compiler, the Developer's Tool Kit (DTK) (OTABASE subset), or a patch that installs a newer version of this library, do not apply this patch. If a new revision of the `libots3` library is already installed on your system, and you install this patch, you will receive the following informational message:

Problem installing:

```
- Tru64_UNIX_V5.1A / Threads Patches
```


Patch 00xxx.00 - Shared libots3 library fix

```
./usr/shlib/libots3.so:  
    is installed by:  
                OTABASE212  
    and cannot be replaced by this patch.
```

This patch will not be installed.

To determine what version of the libots3 library is installed on your system, enter the following command:

```
# what /usr/shlib/libots3.so libots3.so:  
libots3.a      V2.0-094 GEM 27 Feb 2001
```

1.13 Release Note for Tru64 UNIX Patch 252.00

The Essential Services Monitor (ESM) daemon, `esmd`, improves the availability of essential system daemons by automatically restarting them if they terminate. The daemon monitors the Event Manager daemon, `evmd`, and, in a cluster environment, the CAA daemon, `caad`. Restart activity is reported in the `syslog` daemon.log file.

1.14 Release Note for Tru64 UNIX Patch 807.00

This release note updates the `sys_check(8)` reference page.

`syscheck (8)NAME`

`sys_check`, `runsyscheck` - Generates system configuration information and analysis

SYNOPSIS

`/usr/sbin/sys_check [options...]`

OPTIONS

`-all`

Lists all subsystems, including security information and `setld` inventory verification. This option may take a long time to complete.

`-debug`

Outputs debugging information to `stderr` (standard error output).

`-escalate [xx]`

Creates escalation files for reporting problems to your technical support representative. This option produces one file, `TMPDIR/escalate.tar` unless there are crash dump files; if so, it also creates two other files: `TMPDIR/escalate_vmunix.xx.gz` and `TMPDIR/escalate_vmcore.xx.gz`. If you use the `-escalate` option, `sys_check` runs with the `-noquick` option and collects the output in the `escalate.tar` file. Optionally, you can specify a number (`xx`) with the `-escalate` option to define a crash number.

See also the ENVIRONMENT VARIABLES section for information on how you can set the value of `TMPDIR`.

`-evm`

Generates Event Manager (EVM) warnings. When EVM is configured, warnings are posted as EVM events identified by the string `sys.unix.sys_check.warning`. Six levels of priority ranging from 0-500 are used, as follows:

+ 0 - Information only.

+ 100 - Note

- + 200 - Tuning Note
- + 300 - Tuning Suggestion
- + 400 - Operational
- + 500 - Warning

-frame

Produces frame HTML output, which consists of three files: `sys_checkfr.html`, `sys_checktoc.html`, and `sys_check.html` (unless you specify a different file name with the `-name` option). This option cannot be used with the `-nohtml` option. The following options are available for use with the `-frame` option:

-name name

Specifies the name to use for the frame files output. The default name is `sys_check`.

-dir name

Sets the directory for the frames output. Used only with the `-frame` option. The default is the current directory (`.`).

-help or (-h)

Outputs help information.

-nohtml

Produces text output, consisting of one text file, instead of the default HTML output. This option cannot be used with the `-frame` option.

-noquick

Outputs configuration data and the `setld` scan. Excludes security information.

-perf

Outputs only performance data and excludes configuration data. This option takes less time to run than others.

-v Displays the `sys_check` version number.

-warn

Executes only the warning pass. This option takes less time to run than other options.

-nowarn

Executes only the data gathering pass.

DESCRIPTION

The `sys_check` utility is a system census and configuration verification tool that is also used to aid in diagnosing system errors and problems. Use `sys_check` to create an HTML report of your system's configuration (software and hardware). The size of the HTML output that is produced by the `sys_check` utility is usually between .5 MB and 3 MB.

The `sys_check` utility also performs an analysis of operating system parameters and attributes such as those that tune the performance of the system. The report generated by `sys_check` provides warnings if it detects problems with any current settings. Note that while `sys_check` can generate hundreds of useful warnings, it is not a complete and definitive check of the health of your system. The `sys_check` utility should be used in conjunction with event management and system monitoring tools to provide a complete overview and control of system status. Refer to *EVM(5)* for information on event management. Refer to the System Administration guide for information on monitoring your system.

When used as a component of fault diagnosis, `sys_check` can reduce system down time by as much as 50% by providing fast access to critical system data. It is recommended that you run a full check at least once a week to maintain the currency of system data. However, note that some options will take a long time to run and can have an impact on system performance. You should therefore choose your options carefully and run them during off-peak

hours. At a minimum, perform at least one full run (all data and warnings) as a post-configuration task in order to identify configuration problems and establish a configuration baseline. The following table provides guidelines for balancing data needs with performance impact.

Option	Run time	Performance impact	Recommended At
-warn, -perf	Short.	Minimal.	Regular updates, at least weekly
null - no options selected.	Medium, perhaps 15 to 45 minutes depending on processor.	Some likely at peak system use.	Run at least once post-installation and update after major configuration changes. Update your initial baseline and check warnings regularly.
-noquick, -all, -escalate.	Long, perhaps 45 minutes on fast, large systems to hours on low-end systems.	Very likely at peak use.	Use only when troubleshooting a system problem or escalating a problem to your technical support representative.

You can run some `sys_check` options from the SysMan Menu or the `/usr/sbin/sysman -cli` command-line interface. Choose one of the following options from the menu:

```
>- Support and Services
  | Create escalation report [escalation]
  | Create configuration report [config_report]
```

Alternatively, use the `config_report` and `escalation` accelerators from the command line. Note that the `escalation` option should only be used in conjunction with a technical support request.

The `runsyscheck` script will run `sys_check` as a cron task automatically if you do not disable the crontab entry in `/var/spool/cron/crontabs/root`. Check for the presence of an automatically generated log file before you create a new log as it may save time.

When you run the `sys_check` utility without command options, it gathers configuration data excluding the `setld` scan and the security information and displays the configuration and performance data by default. It is recommended that you do this at least once soon after initial system configuration to create a baseline of system configuration, and to consider performing any tuning recommendations.

On the first run, the `sys_check` utility creates a directory named `/var/recovery/sys_check`. On subsequent runs, `sys_check` creates additional directories with a sequential numbering scheme:

- + The previous `sys_check` directory is renamed to `/var/recovery/sys_check.0` while the most recent data (that is, from the current run) is always maintained in `/var/recovery/sys_check`.
- + Previous `sys_check` directories are renamed with an incrementing extension; `/var/recovery/sys_check.0` becomes `/var/recovery/sys_check.1`, and so on, up to `/var/recovery/sys_check.5`.

There is a maximum of seven directories. This feature ensures that you always have up to seven sets of data automatically. Note that if you only perform a full run once, you may want to save the contents of that direc-

tory to a different location.

Depending on what options you choose, the `/var/recovery/sys_check.*` directories will contain the following data:

- + Catastrophic recovery data, such as an `etc` files directory, containing copies of important system files. In this directory, you will find copies of files such as `/etc/group`, `/etc/passwd`, and `/etc/fstab`.
- + Formatted stanza files and shell scripts and that you can optionally use to implement any configuration and tuning recommendations generated by `sys_check` run. You use the `sysconfigdb` command or run the shell scripts to implement the stanza files. See the `sysconfigdb(8)` reference page for more information.

NOTES

You must be root to invoke the `sys_check` utility from the command line; you must be root or have the appropriate privileges through Division of Privileges (DoP) to run Create Configuration Report and Create Escalation Report from the SysMan Menu. The `sys_check` utility does not change any system files.

The `sys_check` utility is updated regularly. You can obtain the latest version of the `sys_check` utility from either of two sources:

- + The most up-to-date version of the `sys_check` kit is located on the `sys_check` tool web site, http://www.tru64unix.compaq.com/sys_check/sys_check.html.
- + You can also obtain `sys_check` from the patch kit, see <http://www.support.compaq.com/patches/>.

You should run only one instance of `sys_check` at a time. The `sys_check` utility prevents the running of multiple instances of itself, provided that the value of the `TMPDIR` environment variable is `/var/tmp`, `/usr/tmp`, `/tmp`, or a common user-defined directory. This avoids possible collisions when an administrator attempts to run `sys_check` while another administrator is already running it. However, no guarantees can be made for the case when two administrators set their `TMPDIR` environment variables to two different user-defined directories (this presumes that one administrator does not choose `/var/tmp`, `/usr/tmp`, or `/tmp`).

The `sys_check` utility does not perform a total system analysis, but it does check for the most common system configuration and operational problems on production systems.

Although the `sys_check` utility gathers firmware and hardware device revision information, it does not validate this data. This must be done by qualified support personnel.

The `sys_check` utility uses other system tools to gather and analyze data. At present, `sys_check` prefers to use `DECEvent`, and you should install and configure `DECEvent` for best results.

If `DECEvent` is not present, the `sys_check` utility issues a warning message as a priority 500 EVM event and attempts to use `uerf` instead. In future releases, Compaq Analyze will also be supported on certain processors.

Note that there are restrictions on using `uerf`, `DECEvent` and Compaq Analyze that apply to:

- + The version of UNIX that you are currently using.
- + The installed version of `sys_check`.
- + The type of processor.

EXIT STATUS

The following exit values are returned:

0 Successful completion.

>0 An error occurred.

LIMITATIONS

DECEvent or Compaq Analyze may not be able to read the binary error log file if old versions of DECEvent are being used or if the binary.errlog file is corrupted. If this problem occurs, install a recent version of DECEvent and, if corrupted, recreate the binary.errlog file.

HSZ controller-specific limitations include the following:

HSZ40 and HSZ50 controllers:

The `sys_check` utility uses a free LUN on each target in order to communicate with HSZ40 and HSZ50 controllers. To avoid data gathering irregularities, always leave LUN 7 free on each HSZ SCSI target for HSZ40 and HSZ50 controllers.

HSZ70, HSZ80 and G80 controllers:

The `sys_check` utility uses a CCL port in order to communicate with HSZ70 controllers. If a CCL port is not available, `sys_check` will use an active LUN. To avoid data gathering irregularities, enable the CCL port for each HSZ70 controller.

The `sys_check` utility attempts to check the NetWorker backup schedule against the `/etc/fstab` file. For some older versions of NetWorker, the `nsradm` command contains a bug that prevents `sys_check` from correctly checking the schedule. In addition, the `sys_check` utility will not correctly validate the NetWorker backup schedule for TruCluster Server.

EXAMPLES

1. The following command creates escalation files that are used to report problems to your technical support organization:
`# sys_check -escalate`
2. The following command outputs configuration and performance information, excluding security information and the `setld` inventory, and provides an analysis of common system configuration and operational problems:
`# sys_check > file.html`
3. The following command outputs all information, including configuration, performance, and security information and a `setld` inventory of the system:
`# sys_check -all > file.html`
4. The following command outputs only performance information:
`# sys_check -perf > file.html`
5. The following command provides HTML output with frames, including configuration and performance information and the `setld` inventory of the system:
`# sys_check -frame -noquick`
6. The following command starts the SysMan Menu `config_report` task from the command line:
`# /usr/sbin/sysman config_report`

Entering this command invokes the SysMan Menu, which prompts you to supply the following optional information:

- + Save to (HTML) - A location to which the HTML report should be saved, which is `/var/adm/hostname_date.html` by default.
- + Export to Web (Default) - Export the HTML report to Insight Manager. Refer to the System Administration manual for information on Insight Manager.
- + Advanced options - This option displays another screen in which you can choose a limited number of run time options. The options

are equivalent to certain command-line options listed in the OPTIONS section.

In this screen, you can also specify an alternate temporary directory other than the default of /var/tmp.

+ Log file - The location of the log file, which is /var/adm/hostname_date.log by default.

7. The following is an example of a stanza file advfs.stanza in /var/recovery/sys_check.*:

```
advfs:
AdvfsCacheMaxPercent=8
```

8. The following is an example of a shell script apply.ksh in /var/recovery/sys_check.*:

```
cd /var/cluster/members/member/recovery/sys_check/
l1ist="advfs.stanza
v1fs.stanza "
for stf in $l1ist; do
print " $stf "
    stanza='print $stf | awk -F . '{print $1 }'
print "/sbin/sysconfigdb -m -f $stf $stanza"
    /sbin/sysconfigdb -m -f $stf $stanza
done
print "The system may need to be rebooted for these
changes to take effect"
```

ENVIRONMENT VARIABLES

The following environment variables affect the execution of the sys_check utility. Normally, you only change these variables under the direction of your technical support representative, as part of a fault diagnosis procedure.

TMPDIR

Specifies a default parent directory for the sys_check working subdirectory, whose name is randomly created; this working subdirectory is removed when sys_check exits. The default value for TMPDIR is /var/tmp.

LOGLINES

Specifies the number of lines of log file text that sys_check includes in the HTML output. The default is 500 lines.

BIGNUMFILE

Specifies the number of files in a directory, above which a directory is considered excessively large. The default is 15 files.

BIGFILE

Specifies the file size, above which a file is considered excessively large. The default is 3072 KB.

VARSIZE

Specifies the minimum amount of free space that sys_check requires in the TMPDIR directory. The default is 15 MB and should not be reduced. The sys_check utility will not run if there is insufficient disk space.

RECOVERY_DIR

Specifies the location for the sys_check recovery data. The default is /var/recovery. The sys_check utility automatically cleans up data from previous command runs. The typical size of the output generated by each sys_check utility run is 400 KB. This data may be useful in recovering from a catastrophic system failure.

ADHOC_DIR

Specifies the location at which sys_check expects to find the text files to include in the HTML output. The default is the /var/adhoc directory.

TOOLS_DIR

Specifies the location at which sys_check expects to find the binaries for the tools that it calls. The default is /usr/sbin.

FILES

`/usr/sbin/sys_check`

Specifies the command path.

Note

This file may be a symbolic link.

`/usr/sbin/*`

Various utilities in this directory are used by `sys_check`.

Note

These files may be symbolic links.

The `sys_check` utility reads many system files.

SEE ALSO

Commands: `dop(8)`, `sysconfigdb(8)`, `sysman_cli(8)`, `sysman_menu(8)`

Miscellaneous: `EVM(5)`, `insight_manager(5)`

Books: *System Administration*, *System Tuning*

1.15 Release Notes for Tru64 UNIX Patch 737.00

This section contains release notes for Patch 737.00.

1.15.1 Updates to `sh`, `csch`, and `ksh`

The updated shells in this kit all implement the following changes when processing shell inline input files:

- File permissions allow only read and write for owner.
- If excessive inline input file name collisions occur, the following error message will be returned:

```
Unable to create temporary file
```

1.15.2 `sh noclobber` Option and `>|` , `>>|` Constructs Added

A `noclobber` option similar to that already available with `csch` and `ksh` has been added to the Bourne shell.

When the `noclobber` option is used (`set -C`), the shell behavior for the redirection operators `>` and `>>` changes as follows:

- For `>` with `noclobber` set, `sh` will return an error rather than overwrite an existing file. If the specified file name is actually a symbolic link, the presence of the symbolic link satisfies the criteria `file exists` whether or not the symbolic link target exists and `sh` returns an error. The `>|` construct will suppress these checks and create the file.
- For `>>` with `noclobber` set, output is appended to the tail of an existing file. If the file name is actually a symbolic link whose target does not exist, `sh` returns an error rather than create the file. The `>>|` construct will suppress these checks and create the file.

1.15.3 `ksh noclobber` Behavior Clarified

For `>` with `noclobber` set, `ksh` will return an error rather than overwrite an existing file. If the specified file name is actually a symbolic link, the presence of the symbolic link satisfies the criteria `file exists` whether or not the symbolic

link target exists and `ksh` returns an error. The `>|` construct will suppress these checks and create the file.

For `>>` with `noclobber` set, output is appended to the tail of an existing file. If the file name is actually a symbolic link to a nonexistent file, `ksh` returns an error. This is a behavior change. Because `ksh` does not have a `>>|` redirection override, create the symbolic link target before accessing it through `>>` if you depend upon appending through a symbolic link.

1.15.4 `csh` `noclobber` Behavior Clarified

For `>` with `noclobber` set, `csh` will return an error rather than overwrite an existing file. If the specified file name is actually a symbolic link, the presence of the symbolic link satisfies the criteria `file exists` whether or not the symbolic link target exists, and `csh` returns an error. The `>|` construct will suppress these checks and create the file.

For `>>` with `noclobber` set, output is appended to the tail of an existing file. If the file does not exist, or the file name is actually a symbolic link whose target does not exist, `csh` returns an error rather than create the file. The `>>|` construct will suppress these checks and create the file.

1.15.5 Updated `mkdir` System Call and Command

This kit reverts the `mkdir` system call, and thus the `mkdir` command, to its Tru64 UNIX Version 4.n behavior with respect to symbolic links. For the unusual case where a symbolic link is used as the very last element of a `mkdir` path, the `mkdir` system call now returns an error than create the target.

If you want `mkdir` to follow the symbolic link you can do so by making the last character of the `mkdir` pathname a slash. For example, if `/var/tmp/foo` is a symbolic link to `/usr/xxx`, which does not exist, then `mkdir("/var/tmp/foo", 0644)` will return an error but `mkdir("var/tmp/foo/", 0644)` will create `/usr/xxx`.

The behavior of `mkdir` can also be controlled systemwide by an addition to the `sysconfig` options for the `vfssubsystem`. The new `sysconfig` option `follow_mkdir_symlinks` defaults to 0, specifying the secure symbolic link behavior. Changing this option to 1, which we strongly discourage, will cause `mkdir` to follow symbolic links.

1.16 Release Note for Tru64 Patch 504.00

This release note contains updates to the `sys_attrs_netrain(5)`, `nifftmt(7)`, `niffconfig(8)`, and `ifconfig(8)` reference pages.

sys_attrs_netrain(5)

`nr_timeout_dead_interface`

The time interval or frequency between successive polls of a dead interface by the NetRAIN interface recovery thread.

Minimum value: 0.5 (seconds)

`nr_timeout_o`

Minimum value: 1.1

`nr_timeout_t`

Minimum value: 0.5

You can specify decimal values (for example, 2.5 or 0.8) for

nr_timeout_dead_interface, nr_timeout_o, and nr_timeout_t. When you reconfigure any of these values by using the sysconfig -r command, they are all validated together. If any value fails validation, all previous (valid) values are restored and EINVAL is returned. Each value must be greater than or equal to its minimum value.

The nr_timeout_o and nr_timeout_t values are validated in conjunction with a third timer value (dt), calculated as $(nr_timeout_t - nr_timeout_o) / 3$. These 3 timer values are validated as described in nifftmt(7).

SEE ALSO

sys_attrs(5), nifftmt(7)

Network Administration: Connections

nifftmt(7)

The time_to_dead field (shown in the EXAMPLES section and in niffconfig -v) is the amount of time that expires between the red alert being raised and the interface being declared dead. It is calculated by the traffic monitor thread as $t2 - t1 - (2 * dt)$.

You can specify the values for t1, dt, and t2 in seconds (if the MIF_MILLISECONDS bit is clear in the flags field), or in milliseconds (if the MIF_MILLISECONDS bit is set). See the EXAMPLES section to see how this is used.

The traffic monitor thread enforces the following restriction between the timing parameters:

$t2 \geq t1 + 2dt$

$t1 \geq 0.5$

$t2 \geq 1.1$

$dt \geq 0.2$

In the preceding restrictions, the values for t1, dt, and t2 are in seconds.

```
#include <stdio.h>
#include <string.h>
#include <sys/types.h>
#include <sys/socket.h>
#include <sys/ioctl.h>
#include <sys/param.h>
#include <net/if.h>
#include <errno.h>

/* these strings map to the "state" enum */
char *state[] = {"INIT", "GREEN", "YELLOW", "ORANGE", "RED", "DEAD"};

/* usage: niff_example tu0 tu1 tu2...
 * must supply the name of at least one
 * network interface
 */
main(int ac, char **av)
{
    int t1 = 20, t2 = 60, dt = 5;
    char **oldav;
    mif_t mif;
    int s;

    oldav = ++av;
    s = socket(AF_INET, SOCK_DGRAM, 0);

    /* tell the traffic monitor to start watching these interfaces */
    while (*av) {
        printf("Adding interface %s to the traffic monitor\n", *av);
        bzero(&mif, sizeof(mif));
        bcopy(*av, &mif.name[0], MIN(strlen(*av)+1, sizeof(mif.name)-1));
    }
}
```

```

mif.t1 = t1;
mif.t2 = t2;
mif.dt = dt;
mif.flags = 0;
if (ioctl(s, SIOCTMTADD, &mif) < 0) {
    perror("couldn't add interface");
    break;
}
++av;
}
av = oldav;

/* get the status of the interfaces - NB will probably always
 * be in the "init" state
 */
while (*av) {
    printf("checking the status of interface %s\n", *av);
    bzero(&mif, sizeof(mif));
    bcopy(*av, &mif.name[0], MIN(strlen(*av)+1, sizeof(mif.name)-1));
    if (ioctl(s, SIOCTMTSTATUS, &mif) < 0) {
        perror("couldn't get status for interface");
        break;
    } else {
        printf("Interface: %05s, state: %s ", mif.name,
            state[mif.current_state]);
        if (mif.flags & MIF_MILLISECONDS)
            printf("Timer values in milliseconds...\n");
        else
            printf("Timer values in seconds...\n");
        printf("t1: %d, dt: %d, t2: %d, time to dead: %d,
            current_interval: %d, next time: %d\n",
            mif.t1, mif.dt, mif.t2, mif.time_to_dead, mif.current_interval,
            mif.next_time);
    }
    ++av;
}
av = oldav;

/* tell the traffic monitor to stop watching */
while (*av) {
    printf("deleting interface %s from the traffic monitor0, *av);
    bzero(&mif, sizeof(mif));
    bcopy(*av, &mif.name[0], MIN(strlen(*av)+1, sizeof(mif.name)-1));
    if (ioctl(s, SIOCTMTREMOVE, &mif) < 0) {
        perror("couldn't remove interface");
    }
    ++av;
}
exit(0);
}

```

niffconfig(8)

SYNOPSIS

```

/usr/sbin/niffconfig [-a] [-m] [-r] [-s] [-u] [-v] [-d dt] [-o t2] [-t t1]
[interface1 interface2...]

```

-d dt

Specifies the time period, in seconds, that the traffic monitor thread uses between reads of the interface counters when it suspects there is a connectivity problem. This number must be smaller than the number given for t1 (see the -t option). The default time period is 5 seconds. If dt is not specified, niffconfig uses the default.

-o t2

Specifies the total number of traffic-free seconds that must elapse before the traffic monitor thread determines that a network interface has failed. This number must be at least the sum of the t1 and two times dt. That is, given the default time period for dt (5 seconds) and t1 (20 seconds), the t2 value must be at least 30 seconds. The default time period for t2 is 60 seconds. If t2 is not specified,

niffconfig uses the default.

- m Modifies the timing parameters of an interface that is already being monitored. Typically, this option is specified along with one or more of -t t1, -d dt, or -o t2 options. If none of these parameters are specified, the default value is used. You cannot specify the -m option with the -a, -s, -r, -u, or -v options.
- t t1
Specifies the time period, in seconds, that the traffic monitor thread delays between reads of the interface counters when the network is running normally. The default time period is 20 seconds. If t1 is not specified, niffconfig uses the default.
- v Displays the status, timer values, and description (verbose mode) of all interfaces currently being monitored to standard out (stdout). See niffmt(7) for a definition of each of the parameters.

Except for the -u and -v options, all niffconfig options require one or more network interfaces to be specified.

You can specify the t1, dt, and t2 timer values as decimal values (for example, 2.6 or 0.8). When setting timer values with the -a or -m options, all three timer values (t1, dt, and t2) are validated as described in niffmt(7). If the validation fails, the operation is cancelled and a message is printed to stdout.

NetRAIN initiates its own internal interface monitoring (using NIFF) when a NetRAIN set is created. NetRAIN monitored interfaces are visible only with the -v option. You cannot use niffconfig to perform any other management operations on the NetRAIN interfaces. To modify the timer values for NetRAIN monitored interfaces, use the ifconfig command.

You can start additional monitoring of an interface that is already being monitored internally for NetRAIN. In that case, the niffconfig -v command will display the two different monitoring structures for the interface. All other niffconfig options will operate only on the non-NetRAIN monitoring structure.

EXAMPLES

5. To display all parameters for all interfaces that are being monitored, including NetRAIN interface monitoring, enter:
niffconfig -v

ifconfig(8)

The `monitor` section should be removed:

The following is added after the second paragraph of the `nrtimers` section:

You can specify decimal values for both the t1 and t2 parameters (for example, 1.5 or 0.8). If you do this, the values are validated similarly to the `nr_timeout_t` and `nr_timeout_o` kernel attributes. See `sys_attrs_netrain(5)` for more information on minimum and maximum NetRAIN timer values.

1.17 Release Note for Tru64 UNIX Patch 846.00

This section contains release notes for Tru64 UNIX Patch 846.00.

1.17.1 Enabling the /dev/poll Function

In order to enable the `/dev/poll` function the special device `poll` must be created manually. The procedure is as follows:

1. Change your directory to `/dev`:

```
# cd /dev
```

2. Execute the MAKEDEV script, found in that directory with either `poll` or `std` as an argument:

```
# MAKEDEV [poll or std]
```

1.17.2 Removal of Version-switched patch

This patch provides a script, `/usr/sbin/evm_versw_undo`, that allows you to remove the EVM patch after the version switch has been thrown by running `clu_upgrade -switch`. This script will set back the version identifiers and request a cluster shutdown and reboot to finish the deletion of the patch. Another rolling upgrade will be required to delete the patch with `dupatch`.

Note

Because the removal of a version-switched patch requires a cluster shutdown, only run this script when you are absolutely sure that this patch is the cause of your problem.

This script must be run by root in multiuser mode after completing the rolling upgrade that installed the patch and before starting another rolling upgrade. The final removal of the patch can only be accomplished by rebooting the system or cluster after this script completes its processing. This script will offer to shut down your system or cluster at the end of its processing. If you choose to wait, it is your responsibility to execute the shutdown of the system or cluster.

Do not forget or wait for an extended period of time before shutting down the cluster. Cluster members that attempt to reboot before the entire cluster is shut down can experience panics or hangs.

1.17.3 New ee Attribute

This patch adds a new `ee` subsystem attribute, `link_check_interval`, that allows the `ee` driver link state polling interval to be tuned for faster failover times when using `ee` devices for Link Aggregation.

The `sys_attrs_ee(5)` reference page is updated as follows:

`link_check_interval`

The interval at which the driver polls the interface link state, in units of hundredths of seconds. Modifying this interval is recommended only when using "ee" devices for Link Aggregation.

Default value: 200 (2 seconds)

Minimum value: 10 (0.1 seconds)

Maximum value: 1000 (10 seconds)

1.17.4 lag(7) Reference Page Update

This patch enables support for network Link Aggregation, or trunking. Link Aggregation can be used to provide increased network bandwidth and availability. Two or more physical Ethernet ports can be combined to create a link aggregation group, which is seen by upper-layer software as a single logical network interface.

See the *Network Administration: Connections* manual for information on configuring link aggregation groups. See `lag(7)` and `lagconfig(8)` for more information about link aggregation.

`lag(7)` `lag(7)`

NAME

lag - Link aggregation (also called trunking) introductory information

DESCRIPTION

Link aggregation, or trunking, enables administrators to combine two or more physical Ethernet Network Interface Cards (NICs) and create a single virtual link. (Upper-layer software sees this link aggregation group as a single virtual interface.) The single virtual link can carry traffic at higher data rates than a single interface because the traffic is distributed across all of the physical ports that make up the link aggregation group.

Using link aggregation provides the following capabilities:

- o Increased network bandwidth - The increase is incremental based on the number and type of ports, or Network Interface Cards (NICs), added to the link aggregation group. See the "Load Sharing" section for more information.
- o Fault tolerance - If a port in a link aggregation group fails, the software detects the failure and reroutes traffic to the other available ports. See the "Fault Tolerance" section for more information.
- o Load sharing - Traffic is distributed across all ports of a link aggregation group. See the "Load Sharing" section for more information.

You can use a link aggregation group virtual interface for the following point-to-point connections: server-to-server and server-to-switch. For server-to-switch connections, the switch must support link aggregation. See your switch documentation for information on configuring your switch.

Link aggregation requires an optional kernel subsystem (lag.mod). You can verify the presence of the link aggregation subsystem by issuing the `sysconfig -s lag` command. If the lag subsystem is not loaded, you can load it using either of the following methods:

- o Dynamically load it using the `sysconfig -c lag` command. This method does not persist across system reboots.
- o Edit the system configuration file, add an options LAG entry to it, and build a new kernel by issuing the `doconfig` command. Then, reboot the system. This method loads the subsystem each time the system reboots.

After the subsystem is loaded, you can configure a link aggregation group.

Link Aggregation Configuration

You can configure link aggregation groups either in multiuser mode or at boot time with the `lagconfig` command. When you configure the group, you can specify a virtual interface number, a key, a distribution algorithm, and a Media Access Control (MAC) address.

After you create a link aggregation group, you can then enable ports (interfaces) for link aggregation. The enabled ports attach to the link aggregation group with the corresponding key. If the port fails in some way, the port detaches from the group and traffic is rerouted to the remaining port or ports.

Any link aggregation configuration done in multiuser mode does not persist across system reboots. If you want link aggregation groups configured at boot time, you must include the appropriate `lagconfig` and `ifconfig` commands in the `/etc/inet.local` file. See the Network Administration: Connections manual for an example.

On platforms where I/O bandwidth may be a limiting factor, you might increase link aggregation performance by distributing the NICs across different portions of the I/O infrastructure (for example, different PCI buses).

Fault Tolerance

The link aggregation subsystem monitors the link state of ports that are enabled for link aggregation. When the link aggregation subsystem detects that a port's link state is down, the subsystem detaches the port from its link aggregation group and redistributes traffic among the remaining ports.

When the link aggregation subsystem detects that the port's link state is up, the subsystem reattaches the port to its link aggregation group. The port then starts handling part of the traffic load again. The amount of time it takes to detect a link state change and fail over depends on the device and driver in use. For DE60x devices using the ee driver in its default configuration, average failover times are on the order of 1 to 2 seconds. To achieve faster failover, reduce the value of the ee subsystem `link_check_interval` attribute. A value of 20 (0.2 seconds) would provide average failover times on the order of 0.1 to 0.2 seconds. For DEGPA devices using the alt driver, average failover times are less than 1 second.

Load Sharing

A link aggregation group performs load sharing of both inbound and outbound traffic. Distribution of inbound packets is determined by the server or switch to which the link aggregation group is connected. When transmitting packets, the system uses a load distribution algorithm to determine on which attached port to transmit the packets. The following load distribution algorithms are supported:

Destination IP Address (dstip)

For IP packets, the port is selected based on a hash of the destination IP address. For non-IP packets, the port is selected based on a hash of the destination MAC address. All traffic addressed to a specific destination IP address uses the same port in the link aggregation group.

This algorithm can utilize the combined bandwidth of a link aggregation group in environments where traffic is destined to a large number of different IP addresses (for example, a web server).

However, this algorithm might not produce the expected bandwidth utilization in environments where the majority of traffic is destined to a single IP address (for example, a private server-to-server interconnect).

Destination MAC address (dstmac)

The port is selected based on a hash of the destination MAC address. All traffic addressed to a specific destination MAC address uses the same port in the link aggregation group.

This algorithm can utilize the combined bandwidth of a link aggregation group in environments where traffic is destined to a large number of different MAC addresses (for example, a server that sends most of its traffic to clients on the same LAN).

However, this algorithm might not produce the expected bandwidth utilization in environments where the majority of traffic is destined to a small number of MAC addresses (for example, a server-to-server interconnect, or a server that sends most of its traffic through a router).

Transport Port number (port)

For TCP or UDP packets originating on the system, the port is selected based on a hash of the source and destination TCP or UDP port numbers. For all other packets, including TCP and UDP packets being forwarded by the system, the Destination IP address (dstip) algorithm is used. All traffic addressed to a specific source+destination port pair uses the same port in the link aggregation group.

This algorithm can utilize the combined bandwidth of a link aggregation group in environments where traffic is destined to a single IP or MAC address, but is exchanged between a number of different TCP or UDP port

number pairs (for example, a server-to-server interconnect).

Round Robin (roundrobin)

The port is selected on a rotating basis.

This algorithm can utilize the combined bandwidth of a link aggregation group in most environments.

However, this algorithm may result in reordering of packets belonging to the same flow (for example, a TCP connection), which in turn may adversely affect performance.

RESTRICTIONS

The following restrictions apply:

- o Supports only DEGPA (alt) and DE60x (ee) network interface cards (NICs).
- o Supports only Ethernet (802.3 CSMA/CD) links.
- o Ports must be operating in full duplex mode.
- o Ports in the same link aggregation group must operate at the same data rate.
- o Ports in a link aggregation group must be attached to the same system, either server-to-server or server-to-switch.

RELATED INFORMATION

Commands: lagconfig(8)

System Attributes: sys_attrs_lag(5)

Files: inet.local(4)

Technical Overview

Network Administration: Connections

1.17.5 lagconfig(8) Reference Page Update

lagconfig(8) lagconfig(8)

NAME

lagconfig - Configures or displays link aggregation groups (or trunk groups)

SYNOPSIS

For creating a link aggregation group, use the following syntax:

```
/usr/sbin/lagconfig -c [attribute,attribute,...]
```

For enabling a port for link aggregation, use the following syntax:

```
/usr/sbin/lagconfig -p port {lag=interface-id | key=value}
```

For deleting a port from a link aggregation group, use the following syntax:

```
/usr/sbin/lagconfig -d port
```

For displaying a link aggregation group, use the following syntax:

```
/usr/sbin/lagconfig -s lag=interface-id
```

OPTIONS

-c Creates a link aggregation group virtual interface. You can specify the following attributes to this option. If you specify more than one attribute, separate them with commas:

lag=interface-id

Specifies the link aggregation group virtual interface name in the form lagn, where n is the unit number (for example, lag1). By default, the next available unit number is assigned to the interface.

key=value

Specifies a value with which to identify the link aggregation group interface. By default, the key value is the next available number. For example, if you previously created a link aggregation group with a key of 4, the next time you create a link aggregation group it is assigned a key of 5.

dist={dstip | dstmac | port | roundrobin}

Specifies the distribution algorithm to be used by the virtual interface for outbound traffic. The software can distribute traffic based on destination IP address (dstip), destination MAC address (dstmac) or transport port number (port), or in a round robin fashion (roundrobin). The default distribution algorithm is dstip. See lag(7) for more information.

macaddr=address

Specifies the Media Access Control (MAC) address to be assigned to the link aggregation group interface. By default, the MAC address of the first link aggregation port (interface) to attach to the link aggregation group is used.

-p port

Enables the specified port (or physical interface) for link aggregation. You must also specify one of the following attributes:

lag=interface-id

Specifies the link aggregation group virtual interface name in the form lagn, where n is the unit number (for example, lag1).

key=value

Specifies the link aggregation group virtual interface to which to add the port by the key assigned to it.

-d port

Deletes the specified port or interface from a link aggregation group.

-s lag=interface-id

Displays the attributes for the specified link aggregation group. The interface-id is in the form lagn, where n is the unit number (for example, lag3).

DESCRIPTION

The lagconfig command allows you to perform the following tasks:

- o Create link aggregation group virtual interfaces.
- o Enable a port (physical interface) for link aggregation.
- o Display attributes for a specified link aggregation group virtual interface.

- o Delete a port from a link aggregation group.

Link aggregation, or trunking, enables administrators to combine one or more physical Ethernet Network Interface Cards (NICs) and create a single virtual link. (Upper-layer software sees this link aggregation group as a single virtual interface.) The single virtual link can carry traffic at higher data rates than a single interface because the traffic is distributed across all of the physical ports that make up the link aggregation group.

If you want to enable a port for link aggregation, you must not configure an IP address on the port, either through the Network Setup Wizard (netconfig) or SysMan. After you enable ports for link aggregation, you issue the `ifconfig up` command to enable the link aggregation group interface. The enabled ports then attach to the link aggregation group that has the same key assigned to it and are available to carry traffic.

If a port fails in some way, the port detaches from the link aggregation group and traffic rerouted to the remaining port or ports. A port also detaches when the system is shut down.

The server or switch at the other end of a link aggregation group must also be configured for link aggregation.

Modifications made with the `lagconfig` command do not persist across reboots of the operating system. To configure the interface or modify the parameters automatically each time the system is booted, edit the `inet.local` file and add the `lagconfig` command and `ifconfig` command entries to it.

Any user can query the status of a link aggregation group; only the superuser can create and modify the configuration of network interfaces.

EXAMPLES

1. To create the link aggregation group virtual interface `lag0` with key value 1 and transport port-based distribution, enter:

```
lagconfig -c lag=lag0,key=1,dist=port
```

2. To add `ee0` and `ee1` to the link aggregation group created in the previous step, enter:

```
lagconfig -p ee0 key=1  
lagconfig -p ee1 key=1
```

Note

Both `ee0` and `ee1` must be DOWN and not have an IP address configured prior to issuing the `lagconfig -p` commands.

3. To display information about the link aggregation group, enter:

```
lagconfig -s lag=lag0  
lag0: Attached Interfaces: ( ee3 ee2 ee1 ee0 )  
key = 1  
Max ports = 8  
dist = port
```

4. To configure an IP address 10.1.2.3 on the link aggregation group virtual interface `lag0` and bring the interface up, enter:

```
ifconfig lag0 10.1.2.3 up
```

DIAGNOSTICS

`lagconfig: subsystem error: Invalid argument`

You attempted to add a port (interface) to a link aggregation group and

the port is UP. Mark the interface DOWN with the `ifconfig` command and try to add the port again.

SEE ALSO

Commands: `netstat(1)`, `ifconfig(8)`, `pfconfig(8)`, `sysconfig(8)`

Interfaces: `lag(7)`

System Attributes: `sys_attrs_lag(5)`

Network Administration: Connections

1.17.6 wol(8) Reference Page Update

This release note contains updates to the `wol(8)` reference page.

`wol(8)`

NAME

`wol` - Send network packet to power on target system (wake-on-LAN)

SYNOPSIS

```
/usr/sbin/wol [nw_interface] hw_address
```

OPTIONS

`nw_interface`

Specifies the network interface to use in making the connection to the target system, for example: `tu1`. This argument is optional.

OPERANDS

`hw_address`

Specifies the hardware network address of the target system, for example: `00-02-56-00-03-29`. This argument is mandatory.

DESCRIPTION

The `wol` utility generates and transmits a network packet to power on a remote system. Before you can use the `wol` utility, you must enable the remote system management wake-on-LAN feature on the target system.

You must specify the target system's hardware address. You may optionally specify the network interface to use in making the connection to the target system. If no network interface is specified, the `wol` utility locates the first configured network interface and prompts you for confirmation.

To enable the wake-on-LAN feature, set the target system's `wol_enable` console variable to `on` and reset the system so that the network controller can read the new state. Use one of the following methods to enable this feature on the target system:

```
+ From the target system's console prompt, enter the following commands:  
>>> set wol_enable on  
>>> init
```

```
+ From the target system's UNIX root prompt, enter the following commands:  
% consvar -s wol_enable on  
set wol_enable = on  
% consvar -a  
Console environment variables saved  
% reboot
```

Use one of the following methods to disable the wake-on-LAN feature:

```
+ From the target system's console prompt, enter the following commands:
```

```
>>> set wol_enable off
>>> init
```

+ From the target system's UNIX root prompt, enter the following commands:

```
% consvar -s wol_enable off
set wol_enable = on
% consvar -a
Console environment variables saved
% reboot
```

Note

You must reset the target system for the new setting to take effect.

RESTRICTIONS

You must be logged in as root or have superuser privileges to use the wol utility.

The wake-on-LAN feature is only available on specific platforms. On platforms that support this feature, additional restrictions may apply. For example, the wake-on-LAN feature may be supported on specific network interface ports only. See your hardware documentation for additional information.

EXIT STATUS

0 (Zero)
Success.

>0 An error occurred.

ERRORS

+ Error detecting default interface

Explanation:

The wol utility cannot automatically detect a default network interface.

User Action:

– Verify that a configured network interface exists on your system.

– Manually specify a configured network interface on the wol command line.

+ Patterns must be specified as hex digits The Magic Packet address must be specified as 00-11-22-33-44-55

Explanation:

The hardware network address entered was in the wrong format. This argument must be in the following format: xx-xx-xx-xx-xx-xx, where x is a hexadecimal character (0 through 9 and A through F, inclusive).

User Action:

Specify the hardware network address correctly.

EXAMPLES

1. The following example shows a simple use of the wol utility, where the host system detects the first configured network interface and prompts for confirmation:

```
# /usr/sbin/wol 00-02-56-00-03-29
No sending device specified, using tu0, continue? (y/n) y
```

2. The following example shows the same use of the wol utility, where the user declines confirmation of the selected network interface:

```
# /usr/sbin/wol 00-02-56-00-03-29
```

No sending device specified, using tu0, continue? (y/n) n
Aborting...

3. The following example explicitly specifies a network interface:

```
# /usr/sbin/wol tu1 00-02-56-00-03-29
```

ENVIRONMENT VARIABLES

wol_enable

Enables or disables the wake-on-LAN feature on the target system. Valid values are on and off.

Note

This is a system console variable, not a UNIX environment variable. The DESCRIPTION section tells you how to enable the wake-on-LAN feature on the target system. You must enable this feature before you use the wol utility.

FILES

/usr/sbin/wol
Wake-on-LAN utility.

SEE ALSO

Commands: consvar(8), halt(8), reboot(8), shutdown(8)

New Hardware Delivery Release Notes and Installation Instructions

System Administration

1.18 Release Note for TruCluster Patch 9.00

This release note explains the relaxed `Cluster Alias: gated` restriction.

Prior to this patch, we required that you use `gated` as a routing daemon for the correct operation of cluster alias routing because the cluster alias subsystem did not coexist gracefully with either the `routed` or `static` routes. This patch provides an `aliasd` daemon that does not depend on having `gated` running in order to function correctly.

The following is a list of features supported by this patch:

- The `gated` and `routed` routing daemons are supported in a cluster. In addition, static routing is supported (no routing daemons are required).

Because `aliasd` is optimized for `gated`, using `gated` remains the default and preferred routing daemon. However, it is no longer mandatory, nor is it the only way to configure routing for a cluster member. For example, you could configure a cluster where all members use static routing, or some members run `routed`, or use a combination of routing daemons and static routes.

However, the existing restriction against using `ogated` still applies; do not use `ogated` as a routing daemon in a cluster.

Note

Cluster members do not have to have identical routing configurations. In general, it is simpler to configure all cluster members identically, but in some instances, an experienced cluster administrator might choose to configure one or more members to perform different routing tasks. For example, one member might have `CLUAMGR_ROUTE_ARGS="nogated"` in its `/etc/rc.config`

file and have a fully populated `/etc/routes` file. Or a member might run with `nogated` and `routed -q`.

- The alias daemon

The alias daemon will handle the failover of cluster alias IP addresses via the cluster interconnect for either dynamic routing or static routing. If an interface fails, `aliasd` reroutes alias traffic to another member of the cluster. As long as the cluster interconnect is working, there is always a way for cluster alias traffic to get in or out of the cluster.

- Interface IP aliases

The `cluamgr` command supports two new `-r` options, `ipalias` and `noipalias`. These options control whether `aliasd` on a member system monitors interface IP aliases. These options let an administrator determine whether a script or `aliasd` manages these interface IP aliases.

When `ipalias` is set, `aliasd` monitors and manages interface IP aliases. When `noipalias` is set, `aliasd` does not monitor or manage IP interface aliases. The default setting is `noipalias`.

Notes

If you use scripts (for example, CAA action scripts) to configure and relocate interface IP aliases for some or all cluster members, run `cluamgr -r noipalias` on those members.

You cannot tell `aliasd` to watch some interface IP aliases on a system but ignore others.

- Multiple interfaces per subnet (for network load balancing)

Although `gated` does not support this configuration, because static routing is supported, an administrator can use static (`nogated`) routing for network load balancing.

By default, the cluster alias subsystem uses `gated`, customized configuration files (`/etc/gated.conf.member<n>`), and RIP to advertise host routes for alias addresses. You can disable this behavior by specifying the `nogated` option to `cluamgr`, either by running the `cluamgr -r nogated` command on a member or by setting `CLUAMGR_ROUTE_ARGS="nogated"` in that member's `/etc/rc.config` file. For example, the network configuration for a member could use `routed`, or `gated` with a site-customized `/etc/gated.conf` file, or static routing.

For a cluster, there are three general routing configuration scenarios:

- The default configuration: `aliasd` controls `gated`.

– Each member has the following in its `/etc/rc.config` file:

```
GATED="yes"
CLUAMGR_ROUTE_ARGS="" # if variable present, set to a null string
```

– If needed, static routes are defined in each member's `/etc/routes` file.

Note

Static routes in `/etc/routes` files are installed before routing daemons are started, and honored by routing daemons.

- Members run `gated`, but the cluster alias and `aliasd` are independent of it. The administrator has total control over `gated` and its configuration file,

`/etc/gated.conf`. This approach is useful for an administrator who wants to enable IP forwarding and configure a member as a full-fledged router.

- Each member that will follow this policy has the following in its `/etc/rc.config` file:

```
GATED="yes"
CLUAMGR_ROUTE_ARGS="nogated"
ROUTER="yes" # if this member will be a full-fledged router
```

- If needed, configure static routes in `/etc/routes`.

- Static routing: one or more cluster members do not run a routing daemon.

- Each member that will use static routing has the following in its `/etc/rc.config` file:

```
GATED="no"
CLUAMGR_ROUTE_ARGS="nogated"
ROUTED="no"
ROUTED_FLAGS=""
```

- Define static routes in that member's `/etc/routes` file.

1.19 Release Note for TruCluster Patch 95.00

When the last member is rolled and right after the version switch is thrown, a script will run which will put CAA on hold and copy the old datastore to the new datastore. CAA will connect to the new datastore when it is available.

The time required to do this depends on the amount of information in the datastore and the speed of each member machine. For 50 resources we have found the datastore conversion itself to only take a few seconds.

To undo this patch, the following command must be run:

```
/usr/sbin/cluster/caa_rollDatastore backward
```

You are prompted to guide the backward conversion process.

One step of this command will prompt you to kill the `caad` daemons on all members. A `caad` daemon may still appear to be running as an uninterruptible sleeping process (state `U` in the `ps` command) after issuing a `kill -9` command. You can safely ignore this and continue with the conversion process as prompted, because `caad` will be killed when the process wakes up.

1.20 Release Note for TruCluster Patch 202.00

This section contains release notes for TruCluster Patch 202.00.

1.20.1 Enablers for EVM

This patch provides enablers for the Compaq SANworks™ Enterprise Volume Manager (EVM) Version 2.0.

1.20.2 Rolling Upgrade Version Switch

This patch uses the rolling upgrade version switch to ensure that all members of the cluster have installed the patch before it is enabled.

Prior to throwing the version switch, you can remove this patch by returning to the rolling upgrade install stage, rerunning `dupatch`, and selecting the Patch Deletion item in the Main Menu.

You can remove this patch after the version switch is thrown, but this requires a shutdown of the entire cluster.

To remove this patch after the version switch is thrown, use the following procedure:

Note

Use this procedure only under the following conditions:

- The rolling upgrade that installed this patch, including the clean stage, has completed.
 - The version switch has been thrown (`clu_upgrade -switch`).
 - A new rolling upgrade is not in progress.
 - All cluster members are up and in multiuser mode.
-

1. Run the `/usr/sbin/evm_ver_sw_undo` command.

When this command completes, it asks whether it should shut down the entire cluster now. The patch removal process is not complete until after the cluster has been shut down and restarted.

If you do not shut down the cluster at this time, you will not be able to shut down and reboot an individual member until the entire cluster has been shut down.

2. After cluster shutdown, boot the cluster to multiuser mode.
3. Rerun the rolling upgrade procedure from the beginning (starting with the setup stage). When you rerun `dupatch`, select the Patch Deletion item in the Main Menu.

For more information about rolling upgrades and removing patches, see the *Patch Kit Installation Instructions*.

1.20.3 Restrictions Removed

The restriction of not supporting multiple filesets from the `cluster_root` domain has been removed. It is now fully supported to have multiple filesets from the `cluster_root` domain to be mounted in a cluster; however, this could slow down the failover of this domain in certain cases and should only be used when necessary.

The restriction of not supporting multiple filesets from a boot partition domain has been removed. It is now fully supported to have multiple filesets from a node's boot partition to be mounted in a cluster; however, when the CFS server node leaves the cluster all filesets mounted from that node's boot partition domain will be force unmounted.

Summary of Base Operating System Patches

This chapter summarizes the base operating system patches included in Patch Kit-0003.

Table 2–1 lists patches that have been updated.

Table 2–2 provides a summary of patches.

Table 2–1: Updated Base Operating System Patches

Patch IDs	Change Summary
Patches 691.00, 693.00, 699.00, 701.00, 705.00, 711.00, 713.00, 718.00, 729.00, 731.00, 733.00, 739.00, 741.00, 749.00, 751.00, 753.00, 755.00, 757.00, 759.00, 761.00, 763.00, 765.00, 767.00, 769.00, 771.00, 773.00, 779.00, 787.00, 789.00, 791.00, 793.00, 795.00, 801.00, 803.00, 809.00, 824.00, 826.00, 828.00, 830.00, 832.00, 840.00	New
Patches 208.00, 429.00, 430.00, 432.00	Superseded by Patch 695.00
Patches 433.00, 434.00, 436.00	Superseded by Patch 697.00
Patches 119.00, 451.00, 3.00, 5.00, 447.00, 449.00, 113.00, 115.00, 706.00, 707.00, 709.00	Superseded by Patch 709.01
Patches 109.00, 110.00, 112.00, 282.00, 284.00, 442.00, 444.00, 714.00	Superseded by Patch 716.00
Patch 471.00	Superseded by Patch 720.00
Patches 472.00, 473.00, 474.00, 475.00, 477.00, 721.00	Superseded by Patch 723.00
Patches 493.00, 495.00	Superseded by Patch 725.00
Patches 496.00, 498.00	Superseded by Patch 727.00
Patches 125.00, 309.00, 460.00, 461.00, 463.00, 734.00, 735.00	Superseded by Patch 737.00
Patches 521.00, 742.00	Superseded by Patch 744.00
Patches 523.00, 745.00	Superseded by Patch 747.00
Patch 584.00	Superseded by Patch 775.00
Patch 586.00	Superseded by Patch 777.00
Patches 80.00, 418.00, 780.00	Superseded by Patch 782.00
Patches 82.00, 420.00, 783.00	Superseded by Patch 785.00
Patches 261.00, 483.00	Superseded by Patch 797.00
Patch 175.00	Superseded by Patch 799.00
Patch 216.00	Superseded by Patch 805.00
Patches 191.00, 2.00, 121.00, 241.00, 243.00, 400.00, 401.00, 402.00, 403.00, 404.00, 405.00, 406.00, 407.00, 408.00, 409.00, 410.00, 412.00, 307.00, 302.00, 162.00, 253.00, 255.00, 90.00, 218.00, 303.00, 305.00, 421.00, 422.00, 423.00, 424.00, 426.00, 681.00, 682.00, 683.00, 684.00, 685.00, 686.00, 687.00, 689.00	Superseded by Patch 807.00
Patch 230.00	Superseded by Patch 811.00
Patch 547.00	Superseded by Patch 813.00

Table 2–1: Updated Base Operating System Patches (cont.)

Patch 549.00	Superseded by Patch 815.00
Patches 165.00, 167.00, 557.00	Superseded by Patch 817.00
Patches 177.00, 559.00, 818.00, 819.00, 820.00	Superseded by Patch 822.00
Patches 67.00, 582.00, 574.00, 576.00, 703	Superseded by Patch 839.00
Patches 541.00, 146.00, 148.00, 126.00, 127.00, 128.00, 129.00, 130.00, 131.00, 132.00, 134.00, 286.00, 6.00, 7.00, 8.00, 9.00, 10.00, 11.00, 12.00, 13.00, 14.00, 15.00, 16.00, 17.00, 18.00, 19.00, 20.00, 21.00, 22.00, 23.00, 24.00, 25.00, 26.00, 27.00, 28.00, 29.00, 30.00, 31.00, 32.00, 33.00, 34.00, 35.00, 36.00, 37.00, 38.00, 39.00, 40.00, 41.00, 42.00, 43.00, 44.00, 45.00, 46.00, 47.00, 48.00, 49.00, 50.00, 51.00, 52.00, 53.00, 54.00, 55.00, 56.00, 57.00, 58.00, 59.00, 60.00, 61.00, 102.00, 63.00, 104.00, 214.00, 236.00, 246.00, 247.00, 248.00, 250.00, 270.00, 271.00, 272.00, 273.00, 274.00, 275.00, 276.00, 277.00, 279.00, 296.00, 298.00, 321.00, 323.00, 325.00, 290.00, 152.00, 93.00, 95.00, 328.00, 329.00, 330.00, 331.00, 332.00, 333.00, 334.00, 335.00, 336.00, 337.00, 338.00, 339.00, 340.00, 341.00, 342.00, 343.00, 344.00, 345.00, 346.00, 347.00, 348.00, 349.00, 350.00, 351.00, 352.00, 353.00, 354.00, 355.00, 356.00, 357.00, 358.00, 359.00, 360.00, 361.00, 362.00, 363.00, 364.00, 365.00, 366.00, 367.00, 368.00, 369.00, 370.00, 371.00, 372.00, 373.00, 374.00, 375.00, 376.00, 377.00, 378.00, 379.00, 380.00, 381.00, 382.00, 383.00, 384.00, 385.00, 386.00, 387.00, 388.00, 389.00, 390.00, 391.00, 392.00, 393.00, 394.00, 395.00, 396.00, 397.00, 399.00, 543.00, 590.00, 592.00, 596.00, 203.00, 594.00, 597.00, 598.00, 599.00, 600.00, 601.00, 602.00, 603.00, 604.00, 605.00, 606.00, 607.00, 608.00, 609.00, 610.00, 611.00, 612.00, 613.00, 614.00, 615.00, 616.00, 617.00, 618.00, 619.00, 620.00, 621.00, 622.00, 623.00, 624.00, 625.00, 626.00, 627.00, 628.00, 629.00, 630.00, 631.00, 632.00, 633.00, 634.00, 635.00, 636.00, 637.00, 638.00, 639.00, 640.00, 641.00, 642.00, 643.00, 644.00, 645.00, 646.00, 647.00, 648.00, 649.00, 650.00, 651.00, 652.00, 653.00, 654.00, 655.00, 656.00, 657.00, 658.00, 659.00, 660.00, 661.00, 662.00, 663.00, 664.00, 665.00, 666.00, 667.00, 668.00, 669.00, 671.00, 672.00, 673.00, 674.00, 675.00, 676.00, 677.00, 678.00, 680.00, 834.00, 835.00, 837.00, 842.00	Superseded by Patch 846.00

Table 2–2: Summary of Base Operating System Patches

Patch IDs	Abstract
Patch 65.00 OSF520-046	Patch: Fix for Compaq C compiler and Compaq driver State: Existing This patch fixes the following problems in the Compaq C compiler and Compiler driver: <ul style="list-style-type: none">• A compiler problem that caused a run-time failure in specific code that involved floating point arguments and varargs.• A problem in the driver that failed to produce an object file for a command such as “file.s -o file.o”.• A problem in the driver that would not allow a command line that contained only the -l<arg> library and no source or object files.• A problem in the driver that failed to produce an object file when no output file was specified on the command line.
Patch 69.00 OSF520-040	Patch: Security (SSRT0743U, SSRT0743U) State: Existing A potential security vulnerability has been discovered where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file or privilege management. We have corrected this potential vulnerability.
Patch 71.00 OSF520CDE-001A	Patch: Security (SSRT1-80U) State: Existing A potential security vulnerability has been discovered where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file or privilege management. We have corrected this potential vulnerability.
Patch 73.00 OSF520CDE-001B	Patch: Security (SSRT1-80U) State: Existing A potential security vulnerability has been discovered where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file or privilege management. We have corrected this potential vulnerability.
Patch 78.00 OSF520X11-007	Patch: Fix for X server hang State: Supersedes patch OSF520X11-006 (76.00) This patch corrects the following: <ul style="list-style-type: none">• Fixes a problem that will cause the X server to hang on rare occasions. Except for the mouse, everything on the desktop appears frozen. Output from the ps command will show the X server using greater than 99% of the CPU time.• Fixes a problem that can cause CDE pop-up menus to appear on the wrong screen when you are running a multihead system with the Panoramix extension enabled.
Patch 84.00 OSF520-143	Patch: Fix for cluster interconnect interface problem State: Existing This patch fixes a problem where shutdown of the network would also shut down the cluster interconnect interface in a LAN cluster.
Patch 86.00 OSF520-054	Patch: Fix for Korn shell hang State: Existing This patch fixes a problem where the Korn shell (ksh) could hang if you pasted a large number of commands to it when it was running in a terminal emulator window (such as an xterm).

Table 2–2: Summary of Base Operating System Patches (cont.)

Patch 88.00 OSF520-022	Patch: Fixes problem with disklabel command State: Existing This patch fixes a problem with the disklabel command. Disklabel was displaying large unsigned values as negative numbers.
Patch 92.00 OSF520-023B	Patch: Security (SSRT1-40U, SSRT1-41U, SSRT1-42U, SSRT1-45U) State: Existing A potential security vulnerability has been discovered where, under certain circumstances, users can clobber temporary files created by shell commands and utilities (for example, under /sbin, /usr/sbin, /usr/bin, and /etc). We have corrected this potential vulnerability.
Patch 97.00 OSF520-001	Patch: Fix for vi editor core dump problem State: Existing This patch fixes a problem where the vi editor core dumps when it finds invalid syntax during a substitute operation.
Patch 106.00 OSF520-026	Patch: Fix for sort command State: Existing This patch corrects the behavior of the sort(1) command, which now checks for duplicates with -c, -u, and -k options.
Patch 108.00 OSF520-015	Patch: Fixes a potential race deadlock State: Existing This patch fixes a potential race deadlock between vclean/ufs_reclaim and quotaon/quotaoff when quota is enabled.
Patch 117.00 OSF520-038	Patch: Fix for evmget command State: Existing This patch fixes a situation in which the evmget command and the event log nightly cleanup operation may fail with an "arg list too long" message.
Patch 123.00 OSF520-056	Patch: Corrects a memory leak in the XTI socket code State: Existing This patch corrects a memory leak in the XTI socket code.
Patch 136.00 OSF520-010A	Patch: Fix for incorrect POSIX 4 message queues behavior State: Existing POSIX 4 message queue behavior was not following the standard and was returning unique message descriptors.
Patch 138.00 OSF520-010B	Patch: Static librt library fix for POSIX 4 message queues State: Existing POSIX 4 message queue behavior was not following the standard and returning unique message descriptors.
Patch 141.00 OSF520X11-005A	Patch: Security (SSRT0638U) State: Supersedes patch OSF520X11-004A (139.00) This patch corrects the following: <ul style="list-style-type: none">• Allows the dxsetacl utility to delete access ACLs.• A potential security vulnerability has been discovered where, under certain circumstances, system integrity may be compromised. This may be in the form of root directory compromise via lpr using X11.
Patch 143.00 OSF520X11-004B	Patch: Allows dxsetacl utility to delete access ACLs State: Existing This patch allows the dxsetacl utility to delete access ACLs.

Table 2–2: Summary of Base Operating System Patches (cont.)

Patch 145.00 OSF520X11-005B	Patch: Security (SSRT0638U) State: Existing A potential security vulnerability has been discovered where, under certain circumstances, system integrity may be compromised. This may be in the form of root directory compromise via lpr using X11.
Patch 154.00 OSF520-061	Patch: Security (SSRT0682U) State: Existing A potential security vulnerability has been discovered where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file or privilege management. We have corrected this potential vulnerability.
Patch 156.00 OSF520DX-004	Patch: Fixes problems which prevented ENVMONd from starting State: Existing This patch fixes problems which prevented ENVMONd from starting
Patch 158.00 OSF520-042	Patch: Fix for Spike postlink optimizer State: Existing This patch fixes a problem where Spike may fail to delete the low instruction of a pair of related instructions, causing it to abort with a run-time error.
Patch 160.00 OSF520-008	Patch: Fix for cp command State: Existing This patch fixes a problem in which cp(1) and cat(1) produce different file sizes when reading from a tape device. The solution changes the I/O buffer size of the cp command from 64 K to 8 K.
Patch 169.00 OSF520-048	Patch: Fixes a problem in latsetup State: Existing This patch fixes a problem in latsetup when the directory /dev/lat is not found.
Patch 171.00 OSF520DX-001	Patch: Fixes a problem in diskconfig State: Existing This fixes a problem in diskconfig where partitions with an offset and size of zero cannot be selected. It also fixes a problem where overlapping partitions cannot be adjusted if the existing partitions are not in alphabetical order.
Patch 173.00 OSF520-076	Patch: Fix for ELSA Gloria Synergy, PS4D10, JIB graphic card State: Existing This patch fixes a problem where, on the ELSA Gloria Synergy, PS4D10, and JIB graphic cards, the cursor position is not being updated properly. The placement of the cursor is one request behind.
Patch 185.00 OSF520-043	Patch: Corrects a problem in the rdist utility State: Existing This patch corrects a problem in the rdist utility which was causing segmentation faults on files with more than one link.
Patch 187.00 OSF520-019	Patch: Fixes a volrecover error State: Existing This patch fixes a volrecover error of "Cannot refetch volume" when volumes exist only in a non-rootdg diskgroup.
Patch 189.00 OSF520-053	Patch: Fix for no rerouting problem on a CFS server State: Existing This patch fixes a problem where pulling the network cable on one node acting as a CFS server in a cluster causes no rerouting to occur.

Table 2–2: Summary of Base Operating System Patches (cont.)

Patch 193.00 OSF520-094	Patch: Support for cleanPR script State: Existing This patch supports the cleanPR script to clear Persistent Reservations on HSV110 device, continues to go through all of devices even if certain errors occur to one or some of devices, and prevent a potential security hole from directly using /tmp directory.
Patch 195.00 OSF520-058	Patch: BPF default packet filter may cause system panic State: This patch corrects a problem which could result in a system panic on close() if the BPF default packet filter is in use.
Patch 210.00 OSF520X11-002	Patch: Fixes problems with X server X Image Extension (XIE) State: Existing This patch fixes problems with the X server X Image Extension (XIE).
Patch 212.00 OSF520-050	Patch: Fixes a problem of the ATM setup script failing State: Existing This patch fixes a problem of the ATM setup script failing when configuring an elan if the lane subsystem is not loaded.
Patch 220.00 OSF520-104	Patch: Corrects a problem with the NIFF daemon State: Existing This patch corrects a problem where the NIFF daemon (niffd) would exit if its connection to the EVM daemon (evmd) failed, as in the case of an EVM daemon restart.
Patch 222.00 OSF520-025	Patch: Fix for mv command State: Existing This patch fixes a problem where the mv command will not perform a move if the inode of the file is the same as the inode of the destination directory, even though the file and directory are on different file systems.
Patch 224.00 OSF520-049	Patch: joinid may fail to clean up its lock files State: Existing The patch fixes a problem where joinid may fail to clean up its lock files in /var/join.
Patch 226.00 OSF520-114A	Patch: Shared libots3 library fix State: Existing This patch fixes the following problems in the /usr/lib/libots3.a and /usr/shlib/libots3.so libraries: <ul style="list-style-type: none">• The max threads clause for the SGI parallel interfaces is being ignored.• An OpenMP thread may hang when reaching a critical region and all other threads are awaiting CVs.
Patch 228.00 OSF520-114B	Patch: Static libots3 library fix State: Existing This patch fixes the following problems in the /usr/lib/libots3.a and /usr/shlib/libots3.so libraries: <ul style="list-style-type: none">• The max threads clause for the SGI parallel interfaces is being ignored.• An OpenMP thread may hang when reaching a critical region and all other threads are awaiting CVs.

Table 2–2: Summary of Base Operating System Patches (cont.)

Patch 234.00 OSF520-124	Patch: Adds support for Persistent Reserve for HSV110 State: Existing This patch is an update to /sbin/scu, the SCSI CAM Utility Program. It adds support for Persistent Reserve for HSV110 as well as the display of 128-bit WWIDS.
Patch 238.00 OSF520DX-002	Patch: Fix for dxsetacl utility State: Existing This patch allows the dxsetacl utility to delete access ACLs.
Patch 245.00 OSF520-173B	Patch: Fixes a problem in the strtod routine State: Existing This patch corrects the following: <ul style="list-style-type: none">• Fixes a problem where strtod() was returning different outputs for the same input.• Fixes a problem where the tan() function was returning the wrong results.
Patch 252.00 OSF520-154	Patch: Adds Essential Services Monitor daemon (esmd) State: Supersedes patch OSF520-099 (75.00) This patch provides enablers for the Compaq Database Utility.
Patch 259.00 OSF520-158	Patch: Removes extraneous header comments State: Existing This patch removes extraneous history edit comments from exported DECThreads header files.
Patch 269.00 OSF520-163	Patch: Improves user control of clu_mibs State: Existing The control of the start and stop of the clu_mibs agent has been moved from /sbin/init.d/clu_max script to /sbin/init.d/snmpd script.
Patch 281.00 OSF520-211	Patch: Fix for NHD kit installations State: Existing During an install of an NHD kit, the version.id file was not properly referenced, causing the install to fail.
Patch 288.00 OSF520-187	Patch: Fix for lpd parent daemon problems State: Existing This patch corrects the following problems: <ul style="list-style-type: none">• Corrects lpd parent daemon problems when EVM is stopped and started.• Slows down event storm from remote host sending bad protocol information.
Patch 295.00 OSF520-169	Patch: Fixes problem of failed open calls to KZPCCs State: Supersedes patch OSF520-195 (293.00) This patch corrects the following: <ul style="list-style-type: none">• Fixes a problem where I/O greater than 4 MB fails to KZPCC devices with error ENODEV.• This patch fixes the problem of failed open calls to KZPCCs under heavy I/O.
Patch 311.00 OSF520DX-012	Patch: Quick Setup erroneously reports daemons do not start State: Existing On some systems, notably DS10, Quick Setup may erroneously report that some daemons did not start. When you then try again, other error messages appear that report duplicate host names.

Table 2–2: Summary of Base Operating System Patches (cont.)

Patch 315.00 OSF520-220B	Patch: Support for Enterprise Volume Manager State: Supersedes patches OSF520-069C (201.00), OSF520-149B (265.00) This patch provides enabler support for the Enterprise Volume Manager.
Patch 319.00 OSF520DX-011	Patch: Security (SSRT1-40U, SSRT1-41U, SSRT1-42U, SSRT1-45U) State: Existing A potential security vulnerability has been discovered where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file access. We have corrected this potential vulnerability.
Patch 327.00 OSF520-167	Patch: Fixes C++ incompatibility State: Existing This patch fixes C++ incompatibility in three files in /usr/include/alpha/hal/ and one file in /usr/include/io/common/.
Patch 414.00 OSF520-230A	Patch: Fixes a problem in stdio.h State: Existing This patch fixes a problem in <stdio.h> where the interface renaming conditionals for fgetpos() and fsetpos() were mismatched. It also fixes a problem in <sys/timeb.h> where the ftime() prototype was not available in the default compilation name space.
Patch 416.00 OSF520-230B	Patch: Fixes a problem in sys/timeb.h State: Existing This patch fixes a problem in <stdio.h> where the interface renaming conditionals for fgetpos() and fsetpos() were mismatched. It also fixes a problem in <sys/timeb.h> where the ftime() prototype was not available in the default compilation name space.
Patch 428.00 OSF520-276B	Patch: Fix for evmwatch termination problem State: Supersedes patches OSF520-103B (164.00), OSF520-159B (257.00) This patch corrects the following: <ul style="list-style-type: none">• Resolves a memory leak and a filtering issue in the Event Manager, and allows the evmwatch utility to reconnect automatically if evmd fails and is restarted.• Fixes a problem in which binary error log (binlog) events posted by the EMX FibreChannel driver and the system console are reported incorrectly by the Event Manager, EVM.• Resolves an issue which can cause an Event Manager (EVM) client or the EVM daemon to core dump under rare circumstances.
Patch 438.00 OSF520DX-016	Patch: Fix for dxproctuner utility State: Existing This patch fixes a problem in dxproctuner where the process information is not displayed when there is a double quote followed by any other character in the command column.

Table 2–2: Summary of Base Operating System Patches (cont.)

Patch 441.00 OSF520X11-019	Patch: Fix for XGetImage function State: Supersedes patches OSF520X11-009 (98.00), OSF520X11-003 (99.00), OSF520X11-001 (101.00), OSF520X11-014 (439.00) This patch corrects the following: <ul style="list-style-type: none">• Provides NHD4 enables for future hardware support of a graphics device.• Fixes the Xserver problem where, when Panoramix is enabled and using CDE, icons from dtfile cannot be seen on other than the left screen while being moved.• Fixes a problem with a Compaq Professional Workstation XP1000 667 MHz system with a PowerStorm 4D20 (PBXGB-CA) graphics card where fonts were sometimes drawn incorrectly.• Fixes a problem where the X Window System XGetImage() function returned erroneous data for displays with a depth greater than 8 when running the Panoramix extension.
Patch 446.00 OSF520DX-024B	Patch: Security (SSRT1-40U, SSRT1-41U, SSRT1-42U, SSRT1-45U) State: Existing A potential security vulnerability has been discovered where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file access. We have corrected this potential vulnerability.
Patch 453.00 OSF520CDE-010	Patch: Fix for dtgreet application State: Existing After installing DCE, enabling SIA would cause a core dump and the greeter window never comes up.
Patch 455.00 OSF520-295	Patch: Fix for lsmsa product State: Existing This patch addresses a problem in the display of disk controller to disk hierarchy by the lsmsa product.
Patch 457.00 OSF520X11-021A	Patch: Fix for broken symbolic links in /usr/lib/X11 State: Existing This patch fixes a problem in Tru64 UNIX V5.1A where three symbolic links in /usr/lib/X11 pointed to nonexistent directories.
Patch 459.00 OSF520X11-021B	Patch: Symbolic links point to nonexistent directories State: Existing This patch fixes a problem in Tru64 UNIX V5.1A where three symbolic links in /usr/lib/X11 pointed to nonexistent directories.
Patch 465.00 OSF520X11-010	Patch: Fix for Elsa Gloria Comet card State: Existing The Elsa Gloria Comet card does not correctly draw nested shaded boxes or anything similar.
Patch 467.00 OSF520X11-013	Patch: Fix for accessx beeping functionality State: Existing Beep does not occur when requested when the toggle keys option is enabled via accessx.
Patch 469.00 OSF520-170	Patch: Fixes a problem in uucp State: Existing This patch fixes a problem in uucp. uucp between two Tru64 UNIX boxes hangs when a uucp failure occurs.

Table 2–2: Summary of Base Operating System Patches (cont.)

Patch 479.00 OSF520-223	Patch: Fix for ACL access problems State: Existing This patch corrects the following: <ul style="list-style-type: none">• If multiple processes attempt to access the same file at the same time and access to the file should be allowed by an ACL on the file, access may be denied instead.• If the ACL on a file is corrupted the corrupted ACL is passed into the kernel causing a variety of problems.
Patch 481.00 OSF520DX-022	Patch: Security (SSRT1-40U, SSRT1-41U, SSRT1-42U, SSRT1-45U) State: Existing A potential security vulnerability has been discovered where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file access. We have corrected this potential vulnerability.
Patch 485.00 OSF520X11-012	Patch: Fix for C++ compile problem State: Existing This patch fixes a C++ compile problem in /usr/include/X11/Xlib.h.
Patch 488.00 OSF520-222A	Patch: Fix for class scheduler State: Supersedes patches OSF520-017A (181.00), OSF520-322 (486.00) This patch corrects the following: <ul style="list-style-type: none">• Fixes a class scheduler semaphore race condition.• Causes the class scheduler to handle rogue programs, changing the class scheduler database semaphore state.• The class scheduler depends on semaphores to protect its database from simultaneous updates. This patch automatically detects if the semaphore no longer exists and allocates a new one, allowing the class scheduler to proceed without interruption.
Patch 490.00 OSF520-222B	Patch: Fix for class scheduler failure State: Supersedes patch OSF520-017B (183.00) This patch corrects the following: <ul style="list-style-type: none">• Fixes a class scheduler semaphore race condition.• The class scheduler depends on semaphores to protect its database from simultaneous updates. This patch automatically detects if the semaphore no longer exists and allocates a new one, allowing the class scheduler to proceed without interruption.
Patch 492.00 OSF520-252	Patch: Fix for verify command State: Existing This patch avoids core dumps in the verify command.
Patch 500.00 OSF520X11-018C	Patch: Security (SSRT0753U, SSRT0752U) State: Existing This patch fixes the libXm.so incompatibility in Tru64 UNIX V5.1A.
Patch 502.00 OSF520DX-013	Patch: Security (SSRT1-40U, SSRT1-41U, SSRT1-42U, SSRT1-45U) State: Existing A potential security vulnerability has been discovered where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file access. We have corrected this potential vulnerability.

Table 2–2: Summary of Base Operating System Patches (cont.)

Patch 504.00 OSF520-171	Patch: Provides faster failover time for NetRAIN State: Supersedes patch OSF520-012 (232.00) This patch corrects the following: <ul style="list-style-type: none">• Fixes a problem in NetRAIN. NetRAIN interface creation now fails if any of the requested standby interfaces do not exist.• In the V5.0 stream NetRAIN failover time has regressed. This patch provides for faster failover time, mainly by permitting timer values of less than 1 second to be configured.
Patch 506.00 OSF520-335	Patch: Fix for rdump command State: Existing The rdump command now dumps data properly onto remote tape devices without receiving the SIGSEGV and dumping core.
Patch 509.00 OSF520-364	Patch: Fix for csh State: Supersedes patch OSF520-182 (507.00) This patch corrects the following: <ul style="list-style-type: none">• If a nonroot user performed an ls(1) with wildcard characters on a directory having permission 700, then it would display the invalid error message, "Glob aborted." Now it displays the correct error message of "Permission denied".• When nonmatch is set and a user performs an ls(1) with one of the patterns as ?, it would not list any matched patterns but return "ls: ? not found". Now it returns that message as well as any matched patterns.• Fixes a problem with the c shell (csh) so that it now correctly recognizes the backslash (\) meta character.
Patch 511.00 OSF520-301	Patch: Fixes alt driver for DEGPA Gigabit Ethernet adapters State: Existing This patch addresses two problems with the alt driver for DEGPA Gigabit Ethernet adapters. These problems affect all Tru64 UNIX systems using alt with vMAC or NetRAIN. <ul style="list-style-type: none">• A fix for vMAC support. Prior to this patch, vMAC has not worked with DEGPA.• A fix to prevent two DEGPA adapters from getting the same MAC address in a NetRAIN configuration.
Patch 513.00 OSF520-333A	Patch: Modifies enablers for Enterprise Volume Manager State: Supersedes patches OSF520-069A (197.00), OSF520-149A (263.00), OSF520-220A (313.00) This patch corrects the following: <ul style="list-style-type: none">• Provides enablers for Enterprise Volume Management.• Modifies enablers for Enterprise Volume Management.
Patch 515.00 OSF520-333B	Patch: Modifies enablers for Enterprise Volume Manager State: Supersedes patches OSF520-069C (201.00), OSF520-149B (265.00), OSF520-220C (317.00) This patch corrects the following: <ul style="list-style-type: none">• Provides enablers for Enterprise Volume Management.• Modifies enablers for Enterprise Volume Management.
Patch 517.00 OSF520-165	Patch: Fix for LSM resynchronization problem State: Existing This patch corrects the problem with a mirrored LSM volume, with dirty region logging (DRL) enabled, still doing a full resynchronization during the first recovery after an unclean shutdown.

Table 2–2: Summary of Base Operating System Patches (cont.)

Patch 519.00 OSF520-155	Patch: Fixes the C++ incompatibility with pwrmgr.h State: Existing This patch fixes the C++ incompatibility of /usr/include/dec/pwrmgr/pwrmgr.h.
Patch 525.00 OSF520-367A	Patch: Security (SSRT0779U) State: Existing A potential security vulnerability has been discovered where, under certain circumstances, SNMP services can stop functioning.
Patch 527.00 OSF520-367B	Patch: Security (SSRT0779U) State: Existing A potential security vulnerability has been discovered where, under certain circumstances, SNMP services can stop functioning.
Patch 529.00 OSF520-174	Patch: Fix for umask permission setting State: Existing This patch fixes a problem where no shell message is displayed when trying to su to a user other than root.
Patch 531.00 OSF520-244	Patch: Fix for KMF caused by malformed IPv4-in-IPv4 packets State: Supersedes patch OSF520-087 (240.00) This patch corrects the following: <ul style="list-style-type: none">• A system configured with the IPTUNNEL kernel option will crash if it receives a corrupted IPv6-in-IPv4 packet, even if the system is not running IPv6. The system will panic with the message "kernel memory fault in ip6ip4_input()"• Fixes a kernel memory fault caused by malformed IPv4-in-IPv4 packets.
Patch 533.00 OSF520-306	Patch: Fix for od command State: Existing This patch fixes a problem in which an invalid character sequence causes the od command to hang or display a partial character.
Patch 535.00 OSF520-251	Patch: Fix for balance utility State: Existing Balance was terminating before balancing the whole domain when the domain was very large (>4 GB).
Patch 537.00 OSF520CDE-007	Patch: Security (SSRT1-40U, SSRT1-41U, SSRT1-42U, SSRT1-48U) State: Supersedes patch OSF520CDE-004 (300.00) A potential security vulnerability has been discovered where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file access. We have corrected this potential vulnerability.
Patch 539.00 OSF520CDE-003	Patch: Security (SSRT0767U) State: Existing A potential security vulnerability has been discovered where, under certain circumstances, system integrity may be compromised. The ttldbserverd contains a potential buffer overflow that may allow unauthorized access. We have corrected this potential vulnerability.
Patch 545.00 OSF520-255	Patch: Fixes EVMs periodic channel monitoring function State: Existing This patch fixes a problem in which the Event Manager's channel monitoring function is temporarily disabled if the evmreload command is run.

Table 2–2: Summary of Base Operating System Patches (cont.)

Patch 551.00 OSF520-181	Patch: Fixes an ATM signaling problem State: Existing This patch fixes an ATM signaling problem.
Patch 553.00 OSF520-317	Patch: EVM daemon fails to find user-defined templates State: Existing This patch resolves a problem with the Event Manager (EVM) where user-defined events are not posted in a semirolled cluster. The Event Manager daemon fails to find user-defined templates in /usr/share/evm/templates/local on nonupgraded nodes in a semirolled cluster.
Patch 555.00 OSF520-219	Patch: Fix for ld linker State: Supersedes patch OSF520-004 (150.00) This patch fixes two problems in the linker (/usr/bin/ld): <ul style="list-style-type: none">• A problem with the datatype of the linker-defined _fpdata symbol.• A problem that causes a linker crash when certain data alignment directives are used in the link.• The linker (/bin/ld) may corrupt the shared object registry file when -update_registry is specified with concurrent links.
Patch 561.00 OSF520X11-016A	Patch: Cut and paste problem with JISX0212 Japanese characters State: Existing This patch fixes a problem with cut and paste of JISX0212 Japanese characters on X Window System applications.
Patch 563.00 OSF520X11-016B	Patch: Fixes JISX0212 Japanese characters problem State: Existing This patch fixes a problem with cut and paste of JISX0212 Japanese characters on X Window System applications.
Patch 565.00 OSF520-324	Patch: Enabler for Compaq Database Utility State: Supersedes patches OSF520-090 (179.00), OSF520-199 (292.00) This patch provides enablers for the Compaq Database Utility.
Patch 567.00 OSF520-210	Patch: Security (SSRT0664U, SSRT0762U) State: Supersedes patches OSF520-045 (204.00), OSF520-068 (206.00) This patch corrects the following: <ul style="list-style-type: none">• This patch corrects a problem with the ftpd daemon which could result in PC ftp clients hanging when transferring some files in ASCII mode.• A potential security vulnerability has been discovered where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file or privilege management. We have corrected this potential vulnerability.• Corrects ftp daemon failure when using a globbing string of several asterisks. Also contains additional corrections for the help command and character drop with the put command.
Patch 569.00 OSF520-294	Patch: CD Mastering Software State: Existing The DS25 system does not include a floppy drive, but has a CD-ROM burner instead. In order to write to this device, CD Mastering Software is required. This patch provides that software. It is made up of mkisofs and cdrecord software.

Table 2–2: Summary of Base Operating System Patches (cont.)

Patch 571.00 OSF520-382	Patch: savecore prematurely terminates crash dump recovery State: Existing This patch corrects a problem where savecore may prematurely terminate crash dump recovery on partitions larger than 4 GB.
Patch 573.00 OSF520DX-014	Patch: Security (SSRT1-40U, SSRT1-41U, SSRT1-42U, SSRT1-45U) State: Existing A potential security vulnerability has been discovered where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file access. We have corrected this potential vulnerability.
Patch 578.00 OSF520-235	Patch: Fix for zdump utility State: Existing This patch corrects the following: <ul style="list-style-type: none">• Fixes a problem in the zdump utility when time zone file names are specified as arguments without leading colons (:).• Fixes a regression in the -v output to display the current time.
Patch 580.00 OSF520X11-020	Patch: Extended Visual Information returns incorrect info State: Existing This patch fixes a problem where the X server's Extended Visual Information (EVI) extension was returning incorrect information.
Patch 588.00 OSF520DX-021	Patch: Fix for NS record syntax in named.local file State: Existing This patch fixes the NS record syntax in a named.local file.
Patch 691.00 OSF520-436B	Patch: Fix for atexit and pthread_prefork handler crashes State: New This patch fixes a problem with atexit() or pthread_atfork() handlers in shared libraries. An application will crash when handlers in shared libraries are called after the libraries are dlclosed and unmapped.
Patch 693.00 OSF520-454B	Patch: Security (SSRT1-40U, SSRT1-41U, SSRT1-42U, SSRT1-45U) State: New A potential security vulnerability has been discovered, where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file access. We have corrected this potential vulnerability.

Table 2–2: Summary of Base Operating System Patches (cont.)

Patch 695.00 OSF520CDE-011A	Patch: Security (SSRT0753U, SSRT0752U, SSRT0788U) State: OSF520CDE-002 (208.00), OSF520CDE-005A (429.00), OSF520CDE-009A (430.00), OSF520CDE-008A (432.00) This patch corrects the following: <ul style="list-style-type: none">• Fixes the problem of palette files not been read from etc/dt/palettes.• Fixes the dtprintinfo memory fault problem with long LANG value.• A potential security vulnerability has been discovered where, under certain circumstances, system integrity may be compromised. This may be in the form of large values of command line arguments. We have corrected this potential vulnerability.• A potential security vulnerability has been discovered , where under certain circumstances, system integrity may be compromised. This may be in the form of large values of ENVIRONMENT variables and command line arguments.• Fixes a potential security vulnerability in CDE Subprocess Control Service(dtspcd). dtspcd has a potential buffer overflow condition which may lead to unauthorized access. We have corrected these potential vulnerability.
Patch 697.00 OSF520CDE-011B	Patch: Security (SSRT0753U, SSRT0752U, SSRT0788U) State: OSF520CDE-005B (433.00), OSF520CDE-009B (434.00), OSF520CDE-008B (436.00) This patch corrects the following: <ul style="list-style-type: none">• Fixes the dtprintinfo memory fault problem with long LANG value.• A potential security vulnerability has been discovered where, under certain circumstances, system integrity may be compromised. This may be in the form of large values of command line arguments. We have corrected this potential vulnerability.• A potential security vulnerability has been discovered , where under certain circumstances, system integrity may be compromised. This may be in the form of large values of ENVIRONMENT variables and command line arguments.• Fixes a potential security vulnerability in CDE Subprocess Control Service(dtspcd). dtspcd has a potential buffer overflow condition which may lead to unauthorized access. We have corrected these potential vulnerability.
Patch 699.00 OSF520-394	Patch: Fix for kdbx State: New This patch fixes a premature termination of the ofile kdbx extension, warning messages in various kdbx extensions, and token length warnings when kdbx is invoked.
Patch 701.00 OSF520-481	Patch: Security (SSRT1-40U, SSRT1-41U, SSRT1-42U, SSRT1-45U) State: New This patch adds the mktemp(1) reference page for the mktemp command.
Patch 705.00 OSF520-485	Patch: Security (SSRT1-40U, SSRT1-41U, SSRT1-42U, SSRT1-45U) State: New This patch updates the mktemp(3) reference page with changed information regarding the mktemp() and mkstemp() routines, and adds information about the mkdtemp() and mkstemp() libc routines.

Table 2–2: Summary of Base Operating System Patches (cont.)

Patch 709.01 OSF520-456	<p>Patch: Fix for autofs and autofs mount daemon</p> <p>State: Supersedes patches OSF520-091 (119.00), OSF520-290 (451.00), OSF520-027 (3.00), OSF520-034 (5.00), OSF520-273 (447.00), OSF520-270 (449.00), OSF520-005 (113.00), OSF520-037 (115.00), OSF520-372 (706.00), OSF520-412 (707.00), OSF520-456 (709.00)</p> <p>This patch corrects the following:</p> <ul style="list-style-type: none">• An AutoFS intercept point for a direct map entry may no longer induce automounts after an error has been detected during a previous automount attempt.• Eliminates error messages concerning property lists seen through certain utilities such as vdump.• AutoFS aut-mounts will now occur when utilities name intercept points defined through indirect map entries.• Fixes a deadlock that will occur in non-cluster systems when direct map entries are served locally.
----------------------------	--

Table 2–2: Summary of Base Operating System Patches (cont.)

Patch 709.00 continued	<ul style="list-style-type: none">• Prevents a core dump from vdump when your message length is greater than MAX_MSG_SIZE. This is a very rare occurrence. The problem was found by code inspection while working on internationalization of messages.• Fixes problems in the vdump command:<ul style="list-style-type: none">– Failed to flag compressed extended attributes records that are split across a vdump BLOCK boundary.– Corrects “Rewinding” message to avoid a segfault with Internationalized messages.• Fixes problems in the vrestore command:<ul style="list-style-type: none">– Fails to properly handle extended attributes records in compressed archives. This results in malloc failures, proplist corruption, program abort, program crashes due to segfault or invalid memory access, and the display of the error message "error setting extended attributes".– Fails to set extended attributes due to confusion over selective restore of the associated file or directory. Also results in display of the error message "error setting extended attributes".– Selective restore of hardlinked files is incomplete when they exist in different directories (fails to create directory for second occurrence of file with same inode number).• Eliminates inefficient behavior by autofs when the top level directory of a direct hierarchical automount map entry cannot be successfully mounted.• Ensures that AutoFS correctly uses the mount options specified in automount map entries with replicated servers.• Fixes a problem where the tar -F (Fasttar) option ignores files named err, but does not ignore files named errs or directories named SCCS and RCS.• Corrects pax/tar/cpio to properly extract explicitly specified files. When an archive contained a file with extended attributes and a different file (occurring later in the archive) was specified to be extracted, improper buffer pointer management resulted in the following display (example uses tar): tar: /dev/nrmt0h : This doesn't look like a tar archive tar: /dev/nrmt0h : Skipping to next file... tar: Memory allocation failed for extended data while reading : Not enough space The directory option was similarly affected. In this case the information for the specified file was not reported• Corrects the behavior of several commands when used in conjunction with file systems that are locally served via AutoFS.• Provides support for wildcards in Linux /etc/exports entries. Both AutoFS and Automount have been so enhanced.• Fixes a problem where aufosd and autofsmount daemon do not properly parse the star (*) in map files.
Patch 711.00 OSF520-339	<p>Patch: Fix for shfragbf</p> <p>State: Existing</p> <p>This patch clarifies the output of shfragbf, an AdvFS utility.</p>

Table 2–2: Summary of Base Operating System Patches (cont.)

Patch 713.00 OSF520-345	Patch: Fix for rcinet script State: Existing This patch prevents the system from hanging when the rcinet script is used by correcting the order in which NetRAIN-related services are started and stopped.
Patch 716.00 OSF520DX-029	Patch: Security (SSRT1-40U, SSRT1-41U, SSRT1-42U, SSRT1-45U) State: Supersedes patches OSF520DX-003 (109.00), OSF520DX-007 (110.00), OSF520DX-006 (112.00), OSF520DX-009 (282.00), OSF520DX-008 (284.00), OSF520DX-015 (442.00), OSF520DX-024A (444.00), OSF520DX-026 (714.00) This patch corrects the following: <ul style="list-style-type: none">• Fixes a problem with the SysMan Station which causes incorrect state information to be displayed after a CPU has been indicated.• Fixes possible deadlock conditions in the SysMan station daemon that might occur at daemon startup or during failover.• Provides enablers for the Compaq Database Utility.• Objects in the Physical File system view do not have correct or updated properties.• SysMan Station can not launch commands on objects where an object attribute is part of the command.• A potential security vulnerability has been discovered where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file access. We have corrected this potential vulnerability.• Fixes a problem where reconfiguration of network interface cards using SysMan makes the old IP address an IP alias. The new IP address replaces the old IP address now.• Fixes a problem where some older, legacy applications were not displaying the appropriate pop-up message because of a change to the way error messages are retrieved. This patch fixes the message catalogs so that the proper message is displayed at the appropriate time.
Patch 718.00 OSF520-426	Patch: Fix for traceroute command State: New This patch corrects a problem where traceroute sometimes failed to provide responses and finish a trace when the destination host name was given on the command line.
Patch 720.00 OSF520-416	Patch: Fix for assembler problems State: Supersedes patch OSF520-239 (471.00) This patch, shipped as Version 3.06.08 of the Tru64 UNIX Assembler, resolves three assembler problems related to the following: <ul style="list-style-type: none">• The generation of an incorrect symbol table which can cause om to fail.• The improper reordering of an instruction which restores the stack pointer when assembling with optimization active.• The generation of a .ident string without a terminating NULL.• When a load instruction specifies a target register and a base register that are the same, the assembler may generate an invalid optimization.

Table 2–2: Summary of Base Operating System Patches (cont.)

Patch 723.00 OSF520DX-025	Patch: Security (SSRT0785U) State: Supersedes patches OSF520DX-010 (472.00), OSF520DX-018 (473.00), OSF520DX-017 (474.00), OSF520DX-019 (475.00), OSF520DX-020 (477.00), OSF520DX-027 (721.00) This patch corrects the following: <ul style="list-style-type: none">• A core dump occurs when <code>/etc/shells</code> is a directory instead of a file.• The hour glass cursor remains after a failure to create a home directory in the process of adding or modifying an account.• Fixes the problem of <code>dxaccounts</code> that names and security attributes of UNIX users are not mapped correctly when they are viewed from the PC Users dialog box.• Fixes the problem that user name entries are replicated in the <code>/etc/group</code> file when modifying users with either <code>dxaccounts</code> or <code>sysman</code> accounts.• Fixes a problem in <code>dxaccounts</code> that can cause certain C2 security values to not be displayed, which could result in unexpected values being saved.• Fixes the problem of <code>useradd</code>, <code>usermod</code>, and <code>dxaccounts</code> ignoring password length restrictions when changing passwords.• Fixes a number of problems with <code>dxaccounts</code> on a system with ASU installed.• A potential security vulnerability has been discovered where, under certain circumstances, system integrity may be compromised. This may be in the form of passwords that have a length outside of the intended range. We have corrected this potential vulnerability.
Patch 725.00 OSF520X11-022A	Patch: Security (SSRT0788U, SSRT0753U, SSRT0752U) State: Supersedes patches OSF520X11-017A (493.00), OSF520X11-018A (495.00) This patch corrects the following: <ul style="list-style-type: none">• A potential security vulnerability has been discovered where, under certain circumstances, system integrity may be compromised. This may be in the form of large values of <code>ENVIRONMENT</code> variables. We have corrected this potential vulnerability.• Fixes the <code>libXm.so</code> incompatibility in Tru64 UNIX V5.1A.• Fixes a problem with the Motif <code>ToggleButton Widget</code> where, in some cases, it may not draw itself correctly.
Patch 727.00 OSF520X11-022B	Patch: Security (SSRT0788U, SSRT0753U, SSRT0752U) State: Supersedes patches OSF520X11-017B (496.00), OSF520X11-018B (498.00) This patch corrects the following: <ul style="list-style-type: none">• A potential security vulnerability has been discovered where, under certain circumstances, system integrity may be compromised. This may be in the form of large values of <code>ENVIRONMENT</code> variables. We have corrected this potential vulnerability.• Fixes the <code>libXm.so</code> incompatibility in Tru64 UNIX V5.1A.• Fixes a problem with the Motif <code>ToggleButton Widget</code> where, in some cases, it may not draw itself correctly.
Patch 729.00 OSF520X11-022B	Patch: Fixes a problem with <code>RLIMIT_DATA</code> process limits State: New This patch fixes a problem with <code>RLIMIT_DATA</code> process limits when running <code>fsck</code> on a large file system.

Table 2–2: Summary of Base Operating System Patches (cont.)

Patch 731.00 OSF520-484	Patch: (SSRT1-40U, SSRT1-41U, SSRT1-42U, SSRT1-45U) State: New This patch adds the safe_open(3) reference page for the safe_open() routine in libc.
Patch 733.00 OSF520-479	Patch: Fix for Memory Channel driver problem State: New Shuts down the Memory Channel device if too many state change interrupts are received.
Patch 737.00 OSF520-445	Patch: Security (SSRT1-40U, SSRT1-41U, SSRT1-42U, SSRT1-45U) State: Supersedes patches OSF520-028 (125.00), OSF520-217 (309.00), OSF520-228 (460.00), OSF520-208 (461.00), OSF520-227 (463.00), OSF520-526 (734.00), OSF520-437 (735.00) This patch corrects the following: <ul style="list-style-type: none">• Fixes a problem in which /usr/bin/ksh hangs for certain scripts that contain wait(1).• A potential security vulnerability has been discovered where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file access. We have corrected this potential vulnerability.• The following changes were made:<ul style="list-style-type: none">– Shell inline input files are more secure.– sh noclobber and new constructs are added.– The mkdir system call is updated.• Corrects a problem in which ksh fails to substitute the tilde (~) character for a user's home directory after an assignment using the number (#) or percent (%) characters has been used.• Fixes a problem with ksh. When a ksh menu is started from within user's .profile, ksh will not stop when the Telnet session is stopped.• Fixes an Asian language processing problem under the Korn shell.• sh will not receive SIGSEGV signal when you run type with file path > 69 characters.• Corrects a problem in which sh was using a high amount of CPU time.• Corrects a problem in which ksh did not clean up the processes associated with a terminal once the window was closed.
Patch 739.00 OSF520-461	Patch: Fixes a problem with os_mibs State: New This patch fixes a problem with os_mibs that could cause the application to consume an excessive amount of CPU time.
Patch 741.00 OSF520-359	Patch: Enhancement to savemeta State: New This patch enhances the capability of savemeta: <ul style="list-style-type: none">• savemeta can now be used in single user mode on a corrupt /usr domain.• All errors now return 1.

Table 2–2: Summary of Base Operating System Patches (cont.)

Patch 744.00 OSF520-466A	Patch: Shared library fix for libaio State: Supersedes patches OSF520-332A (521.00), OSF520-392A (742.00) This patch corrects the following: <ul style="list-style-type: none">• Fixes a rarely seen memory fault in libaio during aio_cancel().• Adds support for NEW_OPEN_MAX_SYSTEM (64K) file descriptors to libaio.• Prevents thread blocking forever when both libaio and libaio_raw are linked into the same image.• Closes an aio_read()/aio_cancel() race condition.
Patch 747.00 OSF520-466B	Patch: Static library fix for libaio State: Supersedes patches OSF520-332B (523.00), OSF520-392B (745.00) This patch corrects the following: <ul style="list-style-type: none">• Fixes a rarely seen memory fault in libaio during aio_cancel().• Adds support for NEW_OPEN_MAX_SYSTEM (64K) file descriptors to libaio.• Prevents thread blocking forever when both libaio and libaio_raw are linked into the same image.• Closes an aio_read()/aio_cancel() race condition.
Patch 749.00 OSF520DX-028	Patch: Modification to secconfig suitlet State: New This patch makes the customize database option available when using secconfig for shadow passwords.
Patch 751.00 OSF520X11-024A	Patch: OSF520X11-024A State: New A potential security vulnerability has been discovered where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file or privilege management. We have corrected this potential vulnerability.
Patch 753.00 OSF520X11-024B	Patch: Security (SSRT0818U) State: New A potential security vulnerability has been discovered where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file or privilege management. We have corrected this potential vulnerability.
Patch 755.00 OSF520X11-024C	Patch: Security (SSRT0818U) State: New A potential security vulnerability has been discovered where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file or privilege management. We have corrected this potential vulnerability.
Patch 757.00 OSF520-362	Patch: Fix for script command State: New This patch corrects a problem in which script would hang upon exit in a dfs configuration.

Table 2–2: Summary of Base Operating System Patches (cont.)

Patch 759.00 OSF520DX-030	Patch: System Mgmt Station detects failing PCI adapters State: New This patch provides the ability for the System Management Station to render PCI adapters with a warning or failed representation when they are in the indicted state. This is in addition to the previous ability to render CPUs that are in the indicted state.
Patch 761.00 OSF520X11-026	Patch: Fixes a problem in the mwm window manager State: New This patch fixes a problem in the mwm window manager where double-click actions are performed on the second button press instead of the second button release. This causes the second button release event to be sent to any underlying window.
Patch 763.00 OSF520-482	Patch: : Security (SSRT1-40U, SSRT1-41U, SSRT1-42U, SSRT1-45U) State: New This patch adds the dirclean(8) reference page for the /usr/sbin/dirclean utility.
Patch 765.00 OSF520-486	Patch: Provides the poll reference page State: New This patch adds the poll(7) reference page for the /dev/poll driver.
Patch 767.00 OSF520-368	Patch: Enhancement to fuser utility State: Existing This patch allows fuser to display the reference option. This option indicates the type of reference made. For example: open, closed, unlinked, or mmaped.
Patch 769.00 OSF520-417	Patch: Fix for su command State: New This patch corrects the behavior of su(1), so that the LOGNAME environment variable is changed to the target user when executed with the - option.
Patch 771.00 OSF520-483	Patch: Lsmsa incorrectly processing passwords State: New This patch fixes a problem where Lsmsa incorrectly processes passwords that are longer than eight characters. Anyone who tries to start the LSM GUI using a password of eight or more characters will be denied access.
Patch 773.00 OSF520-444	Patch: Security (SSRT0795U) State: New A potential security vulnerability has been discovered where, under certain circumstances, system integrity may be compromised. This may be in the form where inetd may block incoming connections when scanned by nmap or other port scanners. We have corrected this potential vulnerability.
Patch 775.00 OSF520CDE-012	Patch: Fixes a problem in the dtwm window manager State: Supersedes patch OSF520CDE-006 (584.00) This patch corrects the following: <ul style="list-style-type: none">• Fixes a memory leak problem in the Window Manager.• Fixes a problem in the dtwm window manager where double-click actions are performed on the second button press instead of the second button release. This causes the second button release event to be sent to any underlying window.

Table 2–2: Summary of Base Operating System Patches (cont.)

Patch 777.00 OSF520-423	Patch: Fixes time zone display problem State: Supersedes patch OSF520-254 (586.00) This patch corrects the following: <ul style="list-style-type: none">• Fixes a problem in which the binlog daemon can coredump if it attempts to recover events from a panic dump file containing invalid event data.• Fixes a time formatting problem when Compaq Analyze is used to display events in time zones with a positive offset from GMT.
Patch 779.00 OSF520X11-023	Patch: Fixes an xfs problem State: New This patch fixes an xfs problem which causes a "QueryGlyphs failed" error in showfont.
Patch 782.00 OSF520-388A	Patch: Added support for DECthreads V3.18-141 State: Supersedes patches OSF520-085A (80.00), OSF520-308A (418.00), OSF520-419A (780.00) This patch corrects the following: <ul style="list-style-type: none">• Installs DECthreads V3.18-133, which fixes problems that may affect threaded programs running on Tru64 UNIX V5.1A. The problems addressed with this patch were discovered during pre-release testing of Tru64 UNIX V5.1A. DECthreads V3.18-133 is the initial support version of the Compaq POSIX Threads Library for Tru64 UNIX V5.1A.• Installs DECthreads V3.18-138, which fixes problems that may affect threaded programs running on Tru64 UNIX V5.1A. This patch specifically addresses a problem that may arise when using recursive mutexes with condition variables.• Installs DECthreads V3.18-144, which is the latest version of the Compaq POSIX Threads library for Tru64 UNIX V5.1A.• Installs DECthreads V3.18-141 which fixes problems that may affect threaded programs running on Tru64 UNIX V5.1A. This patch specifically addresses a problems with the preemption of the symbolic name table() by application code, and the alignment of the Stack Pointer in user created threads.
Patch 785.00 OSF520-388B	Patch: Added support for DECthreads V3.18-141 State: Supersedes patches OSF520-085B (82.00), OSF520-308B (420.00), OSF520-419B (783.00) This patch corrects the following: <ul style="list-style-type: none">• Installs DECthreads V3.18-133, which fixes problems that may affect threaded programs running on Tru64 UNIX V5.1A. The problems addressed with this patch were discovered during pre-release testing of Tru64 UNIX V5.1A. DECthreads V3.18-133 is the initial support version of the Compaq POSIX Threads Library for Tru64 UNIX V5.1A.• Installs DECthreads V3.18-138, which fixes problems that may affect threaded programs running on Tru64 UNIX V5.1A. This patch specifically addresses a problem that may arise when using recursive mutexes with condition variables.• Installs DECthreads V3.18-144, which is the latest version of the Compaq POSIX Threads library for Tru64 UNIX V5.1A.• Installs DECthreads V3.18-141 which fixes problems that may affect threaded programs running on Tru64 UNIX V5.1A. This patch specifically addresses a problems with the preemption of the symbolic name table() by application code, and the alignment of the Stack Pointer in user created threads.

Table 2–2: Summary of Base Operating System Patches (cont.)

Patch 787.00 OSF520-351	Patch: Enhancement for siacfg utility State: New This patch allows nonlocal SIA mechanisms, LDAP for example, to place their mechanism last in the list of mechanisms.
Patch 789.00 OSF520-393	Patch: Fix for ppdof print filter core dump problem State: New This patch corrects a problem where the filter can core dump when the banner jobname contains 132 characters.
Patch 791.00 OSF520-400	Patch: Fix for salvage utility core dump problem State: New This patch fixes a problem with the /sbin/advfs/salvage utility which could cause the utility to core dump.
Patch 793.00 OSF520-343	Patch: Fix for startslip program problem State: New This patch fixes a problem where startslip was not able to extract all the information from the acucap file.
Patch 795.00 OSF520-421	Patch: A timing window can cause a hang in run_usr_cmd State: New A timing window can cause a hang in run_usr_cmd.
Patch 797.00 OSF520-422	Patch: KZPCA devices now work with SCSI devices State: Supersedes patches OSF520-136 (261.00), OSF520-162 (483.00) This patch corrects the following: <ul style="list-style-type: none">• Fixes a panic caused by SCSI bus resets with KZPCA HBAs.• Fixes a kernel memory fault panic after an "ITPSA: itpsa_action - error converting path ID to ITPSA softc structure" message.• Adds the capability for KZPCA devices to work with SCSI devices that only support asynchronous data transfers.
Patch 799.00 OSF520-375	Patch: Fixes several problems with collect utility State: Supersedes patch OSF520-036 (175.00) This patch corrects the following: <ul style="list-style-type: none">• Fixes the Collect's collector (/usr/sbin/collect) to correctly report the network interface load percentage.• Provides fixes for several problems with the Collect utility. The Collect utility has been updated from Version 2.0 to 2.0.3
Patch 801.00 OSF520-431	Patch: Fix for convuser utility State: New Fixes a problem where, if a user was working in enhanced security and then switched to base security, the group and other read privileges would get stripped from /etc/passwd.
Patch 803.00 OSF520-443	Patch: Fixes a problem in access_lines for libmld State: New This patch fixes a problem that may cause the third command and other Atom-based instrumentation tools to fail.
Patch 805.00 OSF520-463	Patch: Fixes a system panic from procfs ioctl user code State: Supersedes patch OSF520-014 (216.00) This patch corrects the following: <ul style="list-style-type: none">• Fixes a kernel memory fault in procfs.mod.• Fixes a system panic from procfs ioctl user code.

Table 2–2: Summary of Base Operating System Patches (cont.)

Patch 807.00 OSF520-480	<p>Patch: Security (SSRT1-40U, SSRT1-41U, SSRT1-42U, SSRT1-45U)</p> <p>State: Supersedes patches OSF520-066 (191.00), OSF520-011 (2.00), OSF520-088 (121.00), OSF520-173A (241.00), OSF520-176 (243.00), OSF520-234 (400.00), OSF520-288 (401.00), OSF520-291 (402.00), OSF520-272 (403.00), OSF520-236 (404.00), OSF520-281 (405.00), OSF520-233 (406.00), OSF520-261 (407.00), OSF520-280 (408.00), OSF520-232 (409.00), OSF520-194 (410.00), OSF520-279 (412.00), OSF520-212 (307.00), OSF520-213 (302.00), OSF520-103A (162.00), OSF520-153 (253.00), OSF520-159A (255.00), OSF520-023A (90.00), OSF520-018 (218.00), OSF520-216 (303.00), OSF520-214 (305.00), OSF520-268 (421.00), OSF520-387 (422.00), OSF520-241 (423.00), OSF520-276A (424.00), OSF520-303 (426.00), OSF520-492 (681.00), OSF520-425 (682.00), OSF520-476 (683.00), OSF520-427 (684.00), OSF520-436A (685.00), OSF520-442 (686.00), OSF520-396 (687.00), OSF520-454A (689.00)</p> <p>This patch corrects the following:</p> <ul style="list-style-type: none">• Fixes a problem where logins appear to be hung on standalone systems with Enhanced Security enabled.• Fixes a regular expression matching problem in multibyte locales.• Fixes the <code>-ignore_all_versions</code> and <code>-ignore_version</code> options for the run-time loader (<code>/sbin/loader</code>).• Fixes a problem where <code>strtod()</code> was returning different outputs for the same input. It also fixes a problem where the <code>tan()</code> function was returning the wrong results.• Eliminates a <code>libc</code> memory leak that occurred when calling <code>dlclose()</code> in applications linked with the thread's run-time environment.• Changes the optional dynamic loader arguments <code>-allocator_range</code> and <code>-allocator</code> to <code>-preallocated_range</code>.• Fixes a problem in <code>mktime()</code> when adjusting for a <code>tm</code> struct containing an invalid <code>tm_isdst</code> (daylight savings time) setting.• A potential security vulnerability has been discovered where, under certain circumstances, system integrity may be compromised. This may be in the form of network programs core dumping. We have corrected this potential vulnerability.• Fixes a segmentation fault problem with long <code>LOCPATH</code> and <code>LANG</code> values.• Fixes a problem in which the RPC TCP server incorrectly tries to write to a socket that has already been closed by a client.• Fixes an application core dump problem when the <code>LANG</code> environment variable is too long.• Fixes a problem with <code>fopen</code>. <code>fopen</code> was returning "file not found" when there was insufficient memory available to allocate the <code>FILE</code> structure. <code>fopen</code> now returns "not enough space" for this case.• Fixes a problem in <code>fread()</code> where excessive I/O was taking place for large amounts of data, causing performance problems. It also addresses a failure in <code>fread()</code> to properly handle data sizes that have representations greater than 32 bits (2^{32} of data).• Fixes a loader core dump that occurs when invoking certain <code>call_shared</code> executables that have been processed by postlink instrumentation tools.• Fixes a problem with <code>strerror</code> where buffers could not be allocated.• Fixes a problem in <code>fwrite()</code> where it was failing when the total number of bytes to be written was larger than 2 GB.• Fixes a regular expression problem with the <code>REG_NEWLINE</code> option of the <code>regex()</code> routine.
----------------------------	---

Table 2–2: Summary of Base Operating System Patches (cont.)

Patch 807.00 continued	<ul style="list-style-type: none">• Fixes a regular expression performance problem as well as two bugs that posed potential regular expression problems for multibyte locales.• A potential security vulnerability has been discovered, where under certain circumstances, system integrity may be compromised. This may be in the form of improper file access. We have corrected this potential vulnerability.• Resolves a memory leak and a filtering issue in the Event Manager, and allows the evmwatch utility to reconnect automatically if evmd fails and is restarted.• Provides enablers for the Compaq Database Utility.• Fixes a problem in which binary error log (binlog) events posted by the EMX FibreChannel driver and the system console are reported incorrectly by the Event Manager, EVM.• A potential security vulnerability has been discovered where, under certain circumstances, users can clobber temporary files created by shell commands and utilities (for example, under /sbin, /usr/sbin, /usr/bin, and /etc). We have corrected this potential vulnerability.• Provides the /usr/sbin/mkstemp program which allows the mechanism to create a secure temporary file.• A potential security vulnerability has been discovered where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file access. We have corrected this potential vulnerability.• Fixes a problem in which the EVM daemon acting as a subscribing client within a cluster will unexpectedly drop the connection to the other EVM daemons in the cluster. This may happen when an EVM client subscribes to events specifying the cluster alias.• Resolves an issue which can cause an Event Manager (EVM) client or the EVM daemon to core dump under rare circumstances.• Fixes the following sys_check problems:<ul style="list-style-type: none">– A potential security vulnerability has been discovered where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file access. We have corrected this potential vulnerability.– The verification of invoking processes' name in CLISCRIPPT failed due to the PARSING of ps output.• Fixes a multi-thread timing window in malloc and free where the list of free chunks could become corrupted, resulting in a segfault.• Fixes a regular expression performance problem in sed.• Fixes a problem with printing long double values.• Fixes a performance degradation in malloc, in applications which perform many mallocs and few frees. With this patch, the performance of malloc is constant regardless of the number of allocated chunks outstanding.• Fixes a problem with atexit() or pthread_atfork() handlers in shared libraries. An application will crash when handlers in shared libraries are called after the libraries are dloaded and unmapped.• Compiled format doprnt code does not handle precision correctly.• A potential security vulnerability has been discovered where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file access. We have corrected this potential vulnerability.
---------------------------	---

Table 2–2: Summary of Base Operating System Patches (cont.)

Patch 809.00 OSF520-527	Patch: Enables correctable error reporting from DTAGII State: New This patch enables correctable error reporting from DTAGII chips on GS320/160/80 1.224Ghz CPU systems.
Patch 811.00 OSF520-465	Patch: Fix for the hwmgr utility State: Supersedes patch OSF520-083 (230.00) This patch fixes three issues with hwmgr: <ul style="list-style-type: none">• An incorrect error message is displayed to the user when using hwmgr to offline a CPU that has only one bound process. The incorrect error message is unable to offline this component and the correct error message should report that there are bound processes on the component.• The path to the scp device is missing when the hwmgr -view devices command is issued.• hwmgr does not show CPU bindings with a tilde (~) whereas the Manage CPU suitlet does when displaying impact information for offlining a CPU. Using the tilde (~) will help distinguish between CPU bindings and RAD bindings. This will also keep the two interfaces consistent for the customer.
Patch 813.00 OSF520X11-025A	Patch: Shared library fix for libXt State: Supersedes patch OSF520X11-015A (547.00) This patch corrects the following: <ul style="list-style-type: none">• Fixes a problem in the X Toolkit library (Xt) which could cause the TeMIP Iconic_map Presentation Module application (mcc_iconic_map) to crash.• Fixes a problem where the definition of the X Toolkit function XtPending() was changed in Tru64 UNIX V5.1A, which caused some applications built on earlier versions of Tru64 UNIX to fail.
Patch 815.00 OSF520X11-025B	Patch: Static library fix for libXt State: Supersedes patch OSF520X11-015B (549.00) This patch corrects the following: <ul style="list-style-type: none">• Fixes a problem in the X Toolkit library (Xt) which could cause the TeMIP Iconic_map Presentation Module application (mcc_iconic_map) to crash.• Fixes a problem where the definition of the X Toolkit function XtPending() was changed in Tru64 UNIX V5.1A, which caused some applications built on earlier versions of Tru64 UNIX to fail.
Patch 817.00 OSF520-446	Patch: Fixes a simple_lock panic when using ATM State: Supersedes patches OSF520-030 (165.00), OSF520-057 (167.00), OSF520-185 (557.00) This patch corrects the following: <ul style="list-style-type: none">• Fixes a kernel memory fault when using ATM.• Corrects a problem which could result in ATM/lane connection requests being dropped.• Fixes a kernel memory fault when using ATM.• Fixes a "simple_lock: time limit exceeded" panic when using ATM.

Table 2–2: Summary of Base Operating System Patches (cont.)

Patch 822.00 OSF520-380	Patch: Fix for fixdmn core dump problem State: Supersedes patches OSF520-065 (177.00), OSF520-260 (559.00), OSF520-370 (818.00), OSF520-397 (819.00), OSF520-391 (820.00) This patch corrects the following: <ul style="list-style-type: none">• Fixes several problems with the fixdmn utility where, under extreme cases, it was possible for fixdmn to core dump or to terminate without fixing the domain.• fixdmn exits prematurely with the message "Can't allocate 0 bytes for group use array" and then instructs user on how to make more memory available, although more memory is not needed.• Allows fixdmn to modify only one page of the transaction log.• Prevents fixdmn from changing file sizes unnecessarily.• Fixes a case where fixdmn would abort when the same mcell was on the DDL more than once.• Allows fixdmn to be run on domains which have been mounted under V5.1B and then moved back to an older OS.• fixdmn could core dump on a rare corruption in the tag file.
Patch 824.00 OSF520-458	Patch: Provides the ckfsec reference page State: New This patch delivers the ckfsec(1) reference page.
Patch 826.00 OSF520-374	Patch: Security (SSRT0794U) State: New A potential security vulnerability has been discovered where, under certain circumstances, system integrity may be compromised. We have corrected this potential vulnerability.
Patch 828.00 OSF520-457	Patch: Security (wc.symlink.002.spautils) State: New A potential security vulnerability has been discovered where, under certain circumstances, system integrity may be compromised. This may be in the form of certain files in world-writable directories. This patch provides the ckfsec utility which can help detect such files.
Patch 830.00 OSF520DX-023	Patch: Incompatibility between Java 1.1.x and Java 2 1.2.x State: New This patch will allow users who are using SysMan tools on the iPAQ, or who have Java 2 as their default Java version, to communicate with Tru 64 UNIX V5.1A systems.
Patch 832.00 OSF520-488	Patch: Update to exportfs reference page State: New This patch updates the exportfs(2) reference page with changed information regarding the exportfsdata structure as a result of increasing a number of file systems that can be NFS mounted from 256 to 1024.

Table 2–2: Summary of Base Operating System Patches (cont.)

Patch 839.00 OSF520-542	Patch: Fix for clsm sync and lsm startup collision problem State: Superseded patches OSF520-105 (67.00), OSF520-226 (582.00), OSF520-331 (574.00), OSF520-318 (576.00), OSF520-447 (703.00) This patch corrects the following: <ul style="list-style-type: none">• Provides enabler for Enterprise Volume Manager product.• Prevents a vold from core dumping when removing a disk from rootdg using voldiskadm or voldg.• Prevents a KMF (kernel memory fault) panic, in voldiskstart(), when an I/O is attempted on an LSM device that is not accessible.• Fixes a situation in which when a cluster member fails, mirrored volumes are left in a state such that recovery is always necessary when members boot, even if no additional recovery should be necessary.• Fixes a clsm sync and lsm startup collision problem.• Fixes a vold core dump when old config db exists.• Fixes cluster node panics on boot if klog does not exist.• Fixes LSM not recognizing third-party disks.• Fixes “unable to create a new diskgroup” when vold is in noloadbalance mode.• Fixes error messages for non-rootdg disks when cluster root is under lsm control.• Fixes problems in LSM’s autoconfiguration feature, as well as some problems in the LSM commands volsave, volrestore, and volclonedg.
Patch 840.00 OSF520-X11-025C	Patch: Definition of X Toolkit function XtPending was changed State: New This patch fixes a problem where the definition of the X Toolkit function XtPending() was changed in Tru64 UNIX V5.1A. This caused some applications built on earlier versions of Tru64 UNIX to fail.
Patch 844.00 OSF520-467B	Patch: AdvFS cannot use volumes greater than 2Tb State: New This patch corrects the following: <ul style="list-style-type: none">• Fixes a bug that prevented AdvFS from working correctly with LSM volumes between 1Tb and 2Tb.• mkfdmn and addvol will now issue a warning if an attempt is made to use an LSM volume greater than 2Tb in size.

Table 2–2: Summary of Base Operating System Patches (cont.)

Patch 846.00	Patch: Security (SSRT1-40U, SSRT1-41U, SSRT1-42U, SSRT1-45U)
OSF520–684	State: Supersedes patches OSF520-304 (541.00), OSF520-118 (146.00), OSF520-071 (148.00), OSF520-078 (126.00), OSF520-077 (127.00), OSF520-126 (128.00), OSF520-007 (129.00), OSF520-115 (130.00), OSF520-121 (131.00), OSF520-009 (132.00), OSF520-074 (134.00), OSF520-207 (286.00), OSF520-097 (6.00), OSF520-081 (7.00), OSF520-116 (8.00), OSF520-044 (9.00), OSF520-020 (10.00), OSF520-021 (11.00), OSF520-138 (12.00), OSF520-089 (13.00), OSF520-128 (14.00), OSF520-075 (15.00), OSF520-031 (16.00), OSF520-142 (17.00), OSF520-141 (18.00), OSF520-039 (19.00), OSF520-127 (20.00), OSF520-033 (21.00), OSF520-024 (22.00), OSF520-120 (23.00), OSF520-029 (24.00), OSF520-051 (25.00), OSF520-052 (26.00), OSF520-131 (27.00), OSF520-055 (28.00), OSF520-059 (29.00), OSF520-130 (30.00), OSF520-098 (31.00), OSF520-129 (32.00), OSF520-035 (33.00), OSF520-064 (34.00), OSF520-109 (35.00), OSF520-100 (36.00), OSF520-101 (37.00), OSF520-062 (38.00), OSF520-106 (39.00), OSF520-117 (40.00), OSF520-125 (41.00), OSF520-063 (42.00), OSF520-016 (43.00), OSF520-096 (44.00), OSF520-092 (45.00), OSF520-112 (46.00), OSF520-108 (47.00), OSF520-133 (48.00), OSF520-137 (49.00), OSF520-067 (50.00), OSF520-032 (51.00), OSF520-086 (52.00), OSF520-111 (53.00), OSF520-147 (54.00), OSF520-080 (55.00), OSF520-047 (56.00), OSF520-073 (57.00), OSF520-107 (58.00), OSF520-002 (59.00), OSF520-060 (60.00), OSF520-151 (61.00), OSF520-113 (102.00), OSF520-070 (63.00), OSF520-110 (104.00), OSF520-123 (214.00), OSF520-093 (236.00), OSF520-150 (246.00), OSF520-156 (247.00), OSF520-172 (248.00), OSF520-168 (250.00), OSF520-183 (270.00), OSF520-192 (271.00), OSF520-203 (272.00), OSF520-196 (273.00), OSF520-186 (274.00), OSF520-191 (275.00), OSF520-204 (276.00), OSF520-201 (277.00), OSF520-205 (279.00), OSF520-221 (296.00), OSF520-215 (298.00), OSF520-247 (321.00), OSF520-284 (323.00), OSF520-313 (325.00), OSF520-189 (290.00), OSF520-119 (152.00), OSF520-079 (93.00), OSF520-084 (95.00), OSF520-274 (328.00), OSF520-305 (329.00), OSF520-248 (330.00), OSF520-237 (331.00), OSF520-299 (332.00), OSF520-293 (333.00), OSF520-309 (334.00), OSF520-316 (335.00), OSF520-275 (336.00), OSF520-277 (337.00), OSF520-250 (338.00), OSF520-193 (339.00), OSF520-206 (340.00), OSF520-242 (341.00), OSF520-320 (342.00), OSF520-188 (343.00), OSF520-209 (344.00), OSF520-337 (345.00), OSF520-177 (346.00), OSF520-307 (347.00), OSF520-256 (348.00), OSF520-330 (349.00), OSF520-285 (350.00), OSF520-132 (351.00), OSF520-267 (352.00), OSF520-152 (353.00), OSF520-271 (354.00), OSF520-298 (355.00), OSF520-297 (356.00), OSF520-245 (357.00), OSF520-328 (358.00), OSF520-184 (359.00), OSF520-240 (360.00), OSF520-262 (361.00), OSF520-180 (362.00), OSF520-190 (363.00), OSF520-259 (364.00), OSF520-356 (365.00), OSF520-157 (366.00), OSF520-198 (367.00), OSF520-258 (368.00), OSF520-197 (369.00), OSF520-315 (370.00), OSF520-325 (371.00), OSF520-360 (372.00), OSF520-286 (373.00), OSF520-140 (374.00), OSF520-266 (375.00), OSF520-326 (376.00), OSF520-342 (377.00), OSF520-278 (378.00), OSF520-327 (379.00), OSF520-296 (380.00), OSF520-314 (381.00), OSF520-166 (382.00), OSF520-302 (383.00), OSF520-202 (384.00), OSF520-310 (385.00), OSF520-263 (386.00), OSF520-264 (387.00), OSF520-257 (388.00), OSF520-319 (389.00), OSF520-311 (390.00), OSF520-253 (391.00), OSF520-323 (392.00), OSF520-329 (393.00), OSF520-287 (394.00), OSF520-238 (395.00), OSF520-145 (396.00), OSF520-231 (397.00), OSF520-265 (399.00), OSF520-334 (543.00), OSF520-338 (590.00), OSF520-418 (592.00), OSF520-501 (596.00), OSF520-102 (203.00), OSF520-407 (594.00), OSF520-378 (597.00), OSF520-487 (598.00), OSF520-385 (599.00), OSF520-478 (600.00), OSF520-429 (601.00), OSF520-369 (602.00), OSF520-464 (603.00), OSF520-432 (604.00), OSF520-358 (605.00), OSF520-401 (606.00), OSF520-414 (607.00), OSF520-347 (608.00), OSF520-502 (609.00), OSF520-468 (610.00), OSF520-474 (611.00), OSF520-377 (612.00),

Table 2–2: Summary of Base Operating System Patches (cont.)

Patch 846.00 continued	OSF520-498 (613.00), OSF520-341 (614.00), OSF520-404 (615.00), OSF520-410 (616.00), OSF520-439 (617.00), OSF520-517 (618.00), OSF520-363 (619.00), OSF520-409 (620.00), OSF520-349 (621.00), OSF520-405 (622.00), OSF520-516 (623.00), OSF520-434 (624.00), OSF520-460 (625.00), OSF520-453 (626.00), OSF520-408 (627.00), OSF520-365 (628.00), OSF520-471 (629.00), OSF520-384 (630.00), OSF520-450 (631.00), OSF520-451 (632.00), OSF520-420 (633.00), OSF520-389 (634.00), OSF520-473 (635.00), OSF520-402 (636.00), OSF520-462 (637.00), OSF520-379 (638.00), OSF520-376 (639.00), OSF520-438 (640.00), OSF520-395 (641.00), OSF520-428 (642.00), OSF520-475 (643.00), OSF520-470 (644.00), OSF520-348 (645.00), OSF520-353 (646.00), OSF520-449 (647.00), OSF520-355 (648.00), OSF520-357 (649.00), OSF520-459 (650.00), OSF520-386 (651.00), OSF520-312 (652.00), OSF520-503 (653.00), OSF520-413 (654.00), OSF520-524 (655.00), OSF520-399 (656.00), OSF520-361 (657.00), OSF520-491 (658.00), OSF520-406 (659.00), OSF520-513 (660.00), OSF520-489 (661.00), OSF520-398 (662.00), OSF520-512 (663.00), OSF520-525 (664.00), OSF520-448 (665.00), OSF520-390 (666.00), OSF520-344 (667.00), OSF520-411 (668.00), OSF520-440 (669.00), OSF520-467 (670.00), OSF520-499 (671.00), OSF520-455 (672.00), OSF520-472 (673.00), OSF520-383 (674.00), OSF520-403 (675.00), OSF520-371 (676.00), OSF520-373 (677.00), OSF520-469 (678.00), OSF520-452 (680.00), OSF520-520 (834.00), OSF520-555 (835.00), OSF520-520 (837.00), OSF520-467A (842.00)
---------------------------	---

This patch corrects the following:

- Fixes “ata_probe: reset failed, sts=0x7f, err=0x7f” errors for IDE disks not connected to the system.
 - Updates the EMX driver to vV2.02 and fixes the following problems:
 - Fixes a panic of “can’t grow probe list”.
 - Fixes a problem of an mcs_lock panic when an adapter experiences a h/w hang condition.
 - Updates the EMX driver to V2.01.
 - Fixes a problem of unexpected tape I/O aborts.
 - Fixes a panic of “can’t grow probe list”.
 - Fixes several kernel memory faults within the driver.
 - Redundant adapter failures no longer panic the system.
 - Corrects a problem of panicking with low memory resources.
 - Corrects stalling I/O during reprobing when a cluster member goes down.
 - Fixes some problems seen with loading and unloading dynamic drivers.
-

Table 2–2: Summary of Base Operating System Patches (cont.)

Patch 846.00 continued	<ul style="list-style-type: none">• Fixes a problem where, when using VX1 graphics module, the mouse cursor disappears when moved along the left and topmost edges.• Fixes a kernel crash dump generation problem which resulted in the wrong page(s) being compressed/written. Without this fix, postmortem debugging may be difficult or impossible.• Fixes a "simple_lock timeout" system panic due to a bug between mcs_unlock and mcs_lock_try on the same CPU.• Provides NHD4 enablers for future hardware support.• Provides a new /usr/sbin/wol command that utilizes the Wake (remotely power) feature for a future platform through the network (LAN).• Provides NHD4 enables for future hardware support of a graphics device.• Fixes a time loss problem seen on DS systems (TSUNAMI) only when using console callbacks. The patch resynchronizes the clock when a time loss is detected.• Fixes a rare panic in the driver for the DE600/DE602 10/100 Ethernet adapter.• Provides NHD4 enablers for future hardware support of a new platform.• Fixes a domain panic pointing to quotaUndo, when a domain has a fileset with a clone, the clone is deleting, and a file in the fileset finds no space available in the domain.• Corrects a problem where the network subsystem sometimes sends a null TCP packet when a connection is reset.• Provides enabler support for Enterprise Volume Manager product.• Fixes a system panic with "malloc_check_checksum: memory pool corruption".• Fixes a problem in which issuing a quot -h command causes a memory fault when the /etc/fstab file contains a mount point that is not mounted.• A potential security vulnerability has been discovered in the kernel where, under certain circumstances, a race condition can occur that could allow a nonroot user to modify any file and possibly gain root access.• Fixes the problem with IPv6 raw socket creations.• Corrects a CFS problem that could cause a panic with the panic string of "CFS_INFS full".• Fixes a problem with erroneous data being returned from the DEVIOCGET ioctl if an error occurs while processing the ioctl.• Fixes a problem in which a TCP socket can continue to receive data with no application running.
---------------------------	---

Table 2–2: Summary of Base Operating System Patches (cont.)

Patch 846.00 continued	<ul style="list-style-type: none">• Fixes a performance problem. The results are large performance increases in configurations where more than 8 tapes are supported on a Fibre Channel (usually behind an MDR or FCTCII).• Allows a single ddr.dbase entry to support a particular SCSI device on both parallel SCSI and FC buses. Previously, SCSI devices connected behind an FCTCII or MDR would not be properly associated with their ddr.dbase entry.• Fixes a panic experienced while task swapping.• Fixes a bug in virtual memory that can cause a kernel memory fault.• Provides NHD4 enablers for future hardware support for an array controller.• Fixes to some problems found with RAID Services that include:<ul style="list-style-type: none">– Raid services not acknowledging presence of CAM RAID device– A hang– The inability to prohibit a user from deleting a logical volume while it is in use– A "malloc_check_checksum: memory pool corruption" system panic• Fixes the following two problems:<ul style="list-style-type: none">– Threads can hang in <code>x_load_inmem_xtnt_map()</code>.– The I/O transfer rate can suddenly drop when writing to a hole in an AdvFS domain, when a volume in that domain becomes full.
---------------------------	--

Table 2–2: Summary of Base Operating System Patches (cont.)

Patch 846.00 continued	<ul style="list-style-type: none">• Fixes the following Virtual Memory problems. The first three are seen on NUMA systems only, and the fourth problem can be seen on any system type:<ul style="list-style-type: none">– A "vm_pg_alloc: page not free" system panic that occurs during process migration.– A "vm_pageout_activate: page already active" system panic that occurs if one thread is unlocking some pages in memory while another thread is migrating them.– Memory inconsistencies caused by the fault path for large shared memory regions prematurely releasing a hold on a page it just locked. This can cause a variety of problems, including user program errors and system panics.– A "simple_lock: time limit exceeded" system panic that occurs if very large (8 MB or larger) System V Shared memory regions are in use.• Fixes a problem with the memory controller attempting to post an EVM event indicating that a particular PFN has been mapped out.• Fixes lock time issues, UBC performance problems, and provides AdvFS and UFS performance improvements in platforms (other than AlphaServer GSxxx) with low memory.• Fixes several bugs related to shared memory (memory that can be accessed by more than one CPU) that could lead to panics, hangs, and performance problems.• Fixes a bug that can cause performance problems for certain applications when the sysconfigtab parameter ipc:sem_broadcast_wakeup is set to 0.• A check for managed address may return an invalid value when called with the address of a gh region not on rad 0.• Fixes a kernel memory fault in msg_rpc_trap.• Fixes a potential problem with lost data after a direct I/O write with a file extension followed quickly by a system crash.• Fixes a crash that occurs when disk controllers are restarted repeatedly.• Fixes a "u_shm_oop_deallocate: reference count mismatch" due to a bug in locking mechanism when gh_chunks are in use.• Provides the I/O barrier code that prevents HSG80 controller crashes (firmware issue).• Corrects the problem of a thread deadlocking against itself under the following conditions:<ul style="list-style-type: none">– Running in a cluster.– Opening (and then closing) a directory that has an index file.– Trying to open the index file through .tags (for example, defragment does that) and by coincidence getting the vnode that pointed to the directory that the index file is attached to.• Fixes a kernel panic with the message "bs_invalidate_rsvd_access_struct: bad access struct".
---------------------------	---

Table 2–2: Summary of Base Operating System Patches (cont.)

Patch 846.00 continued	<ul style="list-style-type: none">• Ensures that DMAPI region information maintains consistency across CFS server and client nodes in the case that an unexpected node failure occurs.• Fixes a problem where additional HSZ70 control ports, /dev/cport/scpN, were created during HSZ70 controller failover operations.• Prevents a crash seen while deleting SCSI devices using hwmgr.• A potential security vulnerability has been discovered where, under certain circumstances, system integrity may be compromised. This could result in a panic with the string: "lock_clear_recursive: recursion not enabled". We have corrected this potential vulnerability.• Fixes a problem where new devices could be created when following the HSZ70 controller failover procedure.• Fixes the problem where reading a clone file that is still in the UBC after an rmvol may panic the system.• Fixes a problem where a variable was used without being initialized, which could lead to a possible kernel memory fault.• Provides AdvFs and VFS support for freezefs and thawfs.• Provides the enabler for Enterprise Volume Manager Version 2.• Corrects several CAM errors including the following:<ul style="list-style-type: none">– Passthru IOCTL fails with EIO (CAM_BUSY) problem.– RESERVATION CONFLICT driver BUSY problem.– Enforces super user-only access for SCSI passthrough.• Enables access to SCSI control ports (/dev/cport/scp??), allowing management of some types of RAID controllers.• Eliminates unintended AutoFS auto-mount storms.• Extraneous "This node removed from cluster" events cause panics of cluster nodes.• Fixes a panic that occurs if DMAPI operations are erroneously executed on an NFS filesystem.• Processes triggering stack growth with anon_rss_enforce set to 2, and exceeding the set resident memory limit, hang, or panic.• Fixes a kernel panic with the messages "xfer_hole_stg: unaligned kernel access" or "xfer_hole_stg: kernel memory fault".
---------------------------	---

Table 2–2: Summary of Base Operating System Patches (cont.)

Patch 846.00 continued	<ul style="list-style-type: none">• Fixes a timing window where flushing data to disk can be incomplete when a system is going down, if more than one thread calls <code>reboot()</code> without first going through <code>shutdown</code>, <code>/sbin/reboot</code>, or <code>/sbin/halt</code>.• Ensures that if an AdvFS file is opened for both <code>O_DIRECTIO</code> and <code>O_APPEND</code>, threads racing to append data to the file will be correctly synchronized, and all data will be appended to the file.• Fixes several direct I/O problems seen when using the AIO interface. The symptoms include a kernel memory fault, and an AIO condition that causes a <code>live_dump</code> to be generated.• Fixes a condition where the <code>smoothsync</code> thread, in attempting to flush dirty buffers for memory-mapped files, would also flush buffers for nonmemory-mapped files. This did not cause any errors, but could cause more I/O than necessary to be done.• Allows POSIX semaphores/msg queues to operate properly on a CFS client.• Fixes the following problems:<ul style="list-style-type: none">– Running <code>verify</code> may panic the system.– A kernel memory fault may occur while attempting to read a log record.• Prevents a race in <code>msfs_umount</code>.• Provides a fix to a deadlock situation that can occur when you invoke the <code>hwmgr -show comp</code> command while the devices on an HSZ70 are changing their names. The devices on an HSZ70 would change their names when you set <code>nofailover</code> or when you set <code>failover</code> on the HSZ70.• Fixes a problem where network interfaces can appear unresponsive to network traffic.• Do not print "path reduced" messages at boot time for devices that still have at least one valid path.• Enables the quick reclaim and deallocation of a <code>vnode</code>.• Under stress conditions where the DMAPI functionality is in use, a panic may occur. A fix is available for this problem.• Fixes a problem where the <code>setgid</code> bit of a directory was not being set when created, if its parent directory has the <code>setgid</code> bit set.• Corrects several problems in kernel routing:<ul style="list-style-type: none">– Fixes a panic when deleting an IP address.– Fixes a panic when performing IP reconfiguration.– Fixes to add interface route on address configuration.• Fixes the panic "<code>ics_unable_to_make_progress: input thread stalled</code>".
---------------------------	---

Table 2–2: Summary of Base Operating System Patches (cont.)

Patch 846.00 continued	<ul style="list-style-type: none">• Addresses three UBC issues:<ul style="list-style-type: none">– Reinstates <code>ubc_maxpercent</code> hardlimit behavior.– Allows the UBC to purge and steal pages under very low free memory conditions during page allocation.– Removes memory mapping for NFS pages being invalidated and freed. Pages were being freed but still mapped the process.• Provides an NFS fix to support the Enterprise Volume Manager product.• Corrects a performance problem where NFS V3 I/O used larger than necessary buffers when writing to locked files resulting in lower throughput.• Provides a script, <code>/usr/sbin/evm_versew_undo</code>, that will allow a user to remove the EVM patch after the version switch has been thrown by running <code>clu_upgrade -switch</code>. This script will set back the version identifiers, request a cluster shutdown, and reboot to finish the deletion of the patch. Another rolling upgrade will be required to delete the patch with <code>dupatch</code>.• Provides an enabler for a version-switched patch.• A SCSI Check Condition with NO SENSE status will now be treated by the disk driver as a condition to retry the I/O.• Fixes a panic that could occur if an illegal argument is passed to UFS mount by a root user.• Fixes a kernel build failure when AdvFs is excluded from the build.• Fixes a problem where the system may be hung or there are poor response times on systems with limited numbers of CPUs.• Fixes an "RDG unwire panic" when running with RDG and GH chunks.• Resolves a problem where duplicate attributes are registered for all CAM devices present in a system. This affects <code>iostat</code> output and any other application that relies on the attribute data.• Adds fixes for additional firmware problems found in the HSx controller.• Fixes the scheduler at high load averages and initial NUMA process placement.• Fixes a <code>rmvol</code> failure that would be seen as an <code>E_PAGE_NOT_MAPPED</code> error when no more space is available for user data migration to another volume in the domain.
---------------------------	---

Table 2–2: Summary of Base Operating System Patches (cont.)

Patch 846.00 continued	<ul style="list-style-type: none">• Fixes the following tape drive problems:<ul style="list-style-type: none">– Tape devices in multipath configurations unexpectedly rewind or go off line. (Multipath means that I/O can reach the device by an alternate data path, such as a redundant controller or bus.) Note that this patch reverts your tape drive configuration to single path mode.– The vdump utility fails to close because the drive goes off line before the dump operation is complete. An error message similar to the following is displayed: vdump: unable to properly close device <dev/tape/tape1_d1>; [5] I/O error• Opening a disk partition sometimes fails when the disk is on shared bus.• Fixes "kernel memory fault" panic on NUMA systems because of corrupt UBC LRU.• Fixes poor interactive response including hanging commands and logins, and random drops in I/O rates when writing many large files.• Fixes a potential problem in which stale data may be returned to an application running on a CFS client when it reads data from a file on a CFS server. Another possible symptom is incomplete flushing of user data when an fsync() is issued or an O_[D]SYNC write is performed.• A potential security vulnerability has been discovered where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file access. We have corrected this potential vulnerability.• New Barrier code will not reserve after a registration if new device or new cluster install.• HSV110 Persistent Reserve with a Reservation conflict SCSI status gets passed off to cam_notify when it should not, resulting in incorrect reservation status.• Addresses a data inconsistency that can occur when a CFS client reads a file using direct I/O that was recently written to.• Fixes a SEL logging problem where panic events were logged as misc events. It also adds new event types that can be logged.• Fixes a problem in which the system could panic while performing CPU hotswap.• A potential security vulnerability has been discovered in networking where, under certain circumstances, a remote system can take over packets destined for another host.• Link Aggregation groups can be successfully created and configured but are unable to successfully transmit and receive packets over the resulting lag interface.• Prevents a potential panic with non-StorageWorks RAID controllers that used the same name for a controller and a disk drive. This conflict was resolved in a prior release but left open the possibility that any attempt to access this disk drive by the kernel could result in a system panic.• Supports a related cluster patch.
---------------------------	---

Table 2–2: Summary of Base Operating System Patches (cont.)

Patch 846.00 continued	<ul style="list-style-type: none">• Removes a panic seen at boot time of the form: panic(cpu 6): u_anon_oop_deallocate: anon_rss_pagelist has pages queued• A potential security vulnerability has been discovered where, under certain circumstances, system integrity may be compromised. This may be in the form of file corruption due to the manner in which setuid/setgid programs core dump. We have corrected this potential vulnerability.• Fixes a kernel memory fault in wait_to_readyq(), or advfs_page_busy(), or potentially other routines which may reference a vm_page, bsBuf, or ioDesc that has been freed prematurely.• Fixes the C++ incompatibility of the following: /usr/include/io/dec/bi/bdareg.h /usr/include/io/dec/bi/buareg.h /usr/include/io/dec/eisa/aceregs.h /usr/include/io/dec/eisa/eisa.h /usr/include/io/dec/fbus/fbusreg.h /usr/include/io/dec/pci/pci.h /usr/include/io/dec/pcmcia/pcmcia.h /usr/include/io/dec/pcmcia/ti1130_reg.h /usr/include/io/dec/tc/sccreg.h /usr/include/io/dec/tc/tc.h /usr/include/io/dec/ws/comet_driver.h /usr/include/io/dec/ws/comet_regs.h /usr/include/io/dec/ws/inputdriver.h /usr/include/io/dec/ws/ws_driver.h• The published cam_logger() interface was modified in V5.1A to accept a hardware ID in its parameter list. This patch restores the cam_logger interface to its published specifications, and introduces the cam_logger3() interface to accept a hardware ID in its parameter list.• Addresses a potential UBC panic which could occur when accessing CFS file systems.• Fixes a problem with vm_faults against anon objects mapped by multiple map entries.• Contains AlphaServer ECC Enhancements for DTAG error logging.• Fixes a problem where decreasing the smoothsync_age does not always have an effect.• Fixes a system panic and/or data inconsistencies caused by changing FIFO parameter pipe-databuf-size while FIFO operations are in flight.• Fixes AdvFS synchronization problems with lingering I/O messages during domain deactivation or rmvol. It also fixes problems caused by certain kmem_debug settings (kmem_debug=0x40, kmem_protected_size=4096) and AdvFS's handling of freed memory.• Fixes and enhances Tru64 UNIX to support Encore realtime software.• Modifies rmvol so that error messages reflect why rmvol fails.• Modifies showfdmn so that showfdmn will not print "Succeeded" on a failure. For example: showfdmn: unable to get info for domain 'domain_used' showfdmn: Successful
---------------------------	---

Table 2–2: Summary of Base Operating System Patches (cont.)

Patch 846.00 continued	<ul style="list-style-type: none">• Fixes a potential CFS deadlock.• Fixes a problem where, when running SSH V2.4.0 and vV2.4.1, users will see a problem executing ls in sftp and when uploading public key using ssh-pubkeymgr.• Fixes SEL logging problem where panic events were logged as misc events. It also adds new event types that can be logged.• Corrects a problem that is encountered when trying to create an Oracle database on an AlphaServer GS system that has a memoryless QBB. Without this patch, direct I/O to an AdvFS file using asynchronous I/O will hang if it is completed on a memoryless QBB.• Corrects problems when running the dd utility on a disk with a label. It would not return errors when expected.• A potential security vulnerability has been discovered where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file or privilege management. We have corrected this potential vulnerability.• Fixes a problem where I/O suspended (hung) in cluster configuration where one or more rad does not have a valid, initialized path.• Fixes a problem that causes bugchecks from applications running DECThreads.• Fixes locking on retry case for multithreaded select/poll. A panic with the following panic string is indicative of this problem: PANIC: "thread_block: simple lock owned"• Fixes a potential problem where system responsiveness may be impacted.• Fixes a Kernel Memory Fault in DMAPi code under cluster stress conditions.• Fixes a calculation leading to poor hash table distribution for NFS client mountpoints in the cluster.• Eliminates unintended AutoFS automounts, in particular those that may result via the execution of any pre-Tru64 UNIX V5.0 df command.• Corrects a problem where multivolume AdvFS V3 domains exhibit I/O errors (not attributable to hardware). The same problem also causes a failed mkfset due to ENO_XTNTS.• Fixes a problem where storage allocation for a file opened for direct I/O could, depending on the write sizes requested, have large extent maps even though the disk was not fragmented. Although the file functioned correctly, performance was reduced by the numerous extent maps. This fix reduces the number of extent maps generated, and subsequently gives better I/O performance on the resulting file.• File permissions inherited from a default ACL may be different than expected in rare cases.• Corrects the problem where the DLI queue stalls when there is no traffic in the TCP/IP or HDLC stacks.• Corrects a problem where clocks on systems could move backwards after subsequent relocations of the root file system using cfsmgr.
---------------------------	---

Table 2–2: Summary of Base Operating System Patches (cont.)

Patch 846.00 continued	<ul style="list-style-type: none">• Two problems are corrected for non-NUMA systems:<ul style="list-style-type: none">– A kernel stack not valid halt on a CPU, which will trigger a PANIC TB_SHOOT ACK TIMEOUT or lock timeout.– A simple lock timeout, or a panic due to holding a simple lock during a context switch.• Corrects an issue seen on NFS clients. The aggressive behavior of client negative lookup cache for concurrent create/lookup was tamed.• Corrects an issue with mmaped() files on an NFS mounted file system. Changes to an mmaped() file were not being immediately seen.• Fixes a problem where the tape changer is only accessible from member that is the drd server for the changer.• Fixes a problem where socket-based applications can hang in soclose().• During file system relocation the system may panic due to a kernel memory fault when a directory larger than 8192 bytes has been deleted while simultaneously being accessed by another thread.• Corrects a kernel memory fault on multiple CPU systems when two or more CPUs find an AdvFS problem at the same time.• Fixes a problem where, after a system crash, on reboot there is a domain panic.• Corrects the problem where attempts to delete psets can hang the system.• Prevents an AdvFS metadata inconsistency in the event of a system crash.• Prevents a possible extent map corruption when multiple volumes are full.• Fixes a problem with multithreaded applications that can cause the application to consume 100% of the CPU usage time.• Fixes a domain panic in a cluster when a file system is mounted on a disk accessed remotely over the cluster interconnect.• Fixes locking problems in vclean().• Fixes the CEH bus/target and lun number when the LUN > 127.
---------------------------	---

Table 2–2: Summary of Base Operating System Patches (cont.)

Patch 846.00 continued	<ul style="list-style-type: none">• Fixes a kernel memory fault when freeing devices.• Corrects problems with USB causing panics under heavily stressed systems.• Corrects a problem with the counters maintained for the NetRAIN virtual interface.• Provides Version 1.02 of the Ciss Driver.• The <code>psrinfo -v</code> command may print an incorrect CPU cache size in a mixed CPU size/speed environment.• Prevents a panic in <code>assert_wait_mesg</code> caused by the posting of an <code>event_wait</code> without clearing a previous request.• Fixes a problem where tape and changer devices on Fibre Channel could occasionally return an incorrect offline status.• Enables the kernel crash dump subsystem to generate a dump after disk driver shutdown has taken place.• Fixes a potential "kernel memory fault" panic in the Virtual Memory subsystem on SMP systems.• Adds hardware support for the DS25, and fixes a minor bug in the ES45 environmental error handling code.• Corrects problems where NFS can deadlock and also corrects an AdvFS problem where EIOs are returned by AdvFS to NFS.• Addresses a kernel memory fault panic in <code>malloc_thread()</code>.• Fixes the predictable TCP Sequence Number.• Addresses a data inconsistency that can occur when a CFS client reads a file that was recently written to.• Supports a related cluster patch to support multiple filesets being mounted from the <code>cluster_root</code> domain.• Fixes a potential deadlock situation when using <code>freezefs</code> on multiple domains while also running <code>addvol</code> (or <code>rmvol</code>).• Fixes numerous problems of accessing deallocated and freed <code>vnodes</code>.• Fixes several problems in Link Aggregation (LAG):<ul style="list-style-type: none">– Cannot modify the <code>ipmtu</code> of a LAG interface.– May attempt to use a link that is down.– Poor performance in server-to-server configurations.• Fixes a situation where a failed open to a device will cause an error so that the device cannot be deleted using <code>hwmgr</code>.• Fixes an incorrect return type in a logging routine that prevented proper operation of the memory troller on a DS20L.• When offlining a processor, a seldom taken code path may attempt to take a complex (sleep or blocking) lock while in interrupt context. Since it is illegal to block in interrupt context, the kernel panics.• Fixes a potential problem with <code>vdf</code> and <code>showfdmn</code>, where they could incorrectly display the message: <code>showfdmn: No such file or directory</code>• Prevents a cluster filesystem-server panic that can occur if a cluster client clears the server cache entries for a file being operated on by <code>defragment</code>, <code>balance</code>, <code>migrate</code>, <code>rmvol</code>, or <code>mssh</code>.
---------------------------	--

Table 2–2: Summary of Base Operating System Patches (cont.)

Patch 846.00 continued	<ul style="list-style-type: none">• Fixes several problems found in the KZPEA driver that could result in hung I/O, pending I/O not being cleared on a reset, panics seen when aborting I/O, and a hard error returned to applications on opens during reset conditions.• A potential security vulnerability has been discovered where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file access. We have corrected this potential vulnerability.• Contains changes to the evm_versw_undo script to fix no-roll installation and deletion of the EVM version-switched patch.• Fixes a problem with the logging of MUNSA reject status messages to the console during boot which could cause a system to boot extremely slow.• Fixes a kernel memory fault from sth_close_fifo() caused by a NULL pointer.• Fixes an ISO9660 file system size limitation of 2.1 GB and provides full capacity access to DVDROM media.• Cluster unlinked files are now handled properly during a relocation.• Fixes a deadlock problem when deleting devices while the system disk is in error recovery.• Recovers the POSIX semantics for accessing "." entry.• Closes a race condition between VFS and UFS layer code that causes panic while periodic sync mechanism flushes dirty buffers out to disks.• Fixes performance shortcomings in NXM thread replacement.• Reduces the number of inputs/outputs to the disk, which reduces the number of audible disk ticks. <p>Fixes a problem for NFS mounts using the -o proplist option, where disk space is not being freed when files are deleted.</p> <ul style="list-style-type: none">• Fixes a kernel memory fault in u_seg_global_destroy.• Fixes a crash when an AdvFS filesystem reports I/O errors and enters into a domain panic state. AdvFS's error cleanup would panic on an invalid pointer and report an "invalid memory read access from kernel mode" panic message.• Fixes an ISO9660 file system size limitation of 2.1GB and provides full capacity access to DVDROM media.• Fixes a problem which can prevent certain tape applications from recovering paths to devices that have failed.• Prevents a panic from occurring while trying to mount an AdvFS domain. The panic would only appear when the mount command was about to fail.• Fixes a problem with the logging of MUNSA reject status messages to the console during boot which could cause a system to boot extremely slowly. It also corrects a problem with reservation conflicts in the TUR recovery loops.• Fixes two NFS kernel memory fault panics due to bad NFS server data.• Corrects "u_anon_free: page busy" panics.• Corrects a problem where an I/O can fail back to the application when a HSV110 V2 path failover is performed. This is incorrect behavior; the application should not have the I/O returned with a failure status when a path failover occurs.
---------------------------	---

Table 2–2: Summary of Base Operating System Patches (cont.)

Patch 846.00 continued	<ul style="list-style-type: none">• Corrects a problem where a tape open with no tape present in the drive can take as long as six minutes to fail.• Fixes a VFS namecache race condition where both positive and negative namecache entries can exist.• Corrects a problem where low UBC memory conditions cause hang in AdvFS.• Corrects a problem in cluster backups of global root directories or backup of different system disks in a cluster.• Fixes an AdvFS alignment fault panic. The panic is caused by inconsistent AdvFS metadata in a directory. In particular, the directory's entry size is an unaligned value.• If an I/O fails and it may be helped by an AdvFS-initiated retry, a message will be written to the console providing information on how to retry.• Fixes a problem where an internal value related to the maximum I/O size for a device was being calculated incorrectly.• Provides support for the DEGXA Gigabit Ethernet device, including the ES25 onboard 10/100/1000 Ethernet port.• Adds the Reserve/Release/Path_Lock feature to tape and changer drivers.• Corrects a problem of writing erroneous data to the binary error log file and provides missing header definitions for error interpretation.• Contains a variety of domain panic fixes that better capture, explain, and handle domain panics.• Fixes a cluster-as-NFS-server chown() problem.• Fixes a problem where the system panics while running applications performing open of RAID device, and the faulting routine is control_port_open.• Fixes a problem where cluster filesystem I/O and AdvFS domain access causes processes to hang.• Fixes a situation in which a system running cdrecord software experiences I/O timeout errors when writing CDs.• Corrects a silent infinite loop in vdump by correcting the AdvFS system call OP_GET_BKUP_XTNT_MAP. The call will now return the valid xtntCnt when it fails due to E_NOT_ENOUGH_XTNTS.• Fixes a problem where, in certain cases, large files (~30Gb) suffered extreme fragmentation.
---------------------------	--

Table 2–2: Summary of Base Operating System Patches (cont.)

Patch 846.00 continued	<ul style="list-style-type: none">• Fixes a kernel panic of "bfs_alloc: kernel memory fault".• Corrects a problem which had resulted in broadcast or multicast packets being processed multiple times on behalf of a NetRAIN device, once for each backup interface.• Fixes a problem that caused the 4.3BSD socket interface to return incorrect values for IOCTL calls accessing IP alias address information.• Supports a cluster patch which corrects a performance issue seen when multiple threads/processes simultaneously access the same file on an SMP (>1 CPU) system.• Fixes a problem where, when entering the hwmgr -view devices command on a member in a cluster, the device name would not be updated and would be listed as unknown.• When adding a CDROM or floppy disk after boot, only A and C device special files would get created. In order to create all the device special files, the user would have to enter dsfmgr -K or reboot. This no longer needs to be done by the user.• Fixes heap and stack limitations in the older OS versions required for SAP.• Prevents several possible system panics and an AdvFS deadlock.• Fixes a problem that allowed an application with superuser privileges to cause a system panic when attempting to delete a nonexistent connection; for example, when the program tcpkill runs while stopping ASU.• Fixes an AdvFS AIO read timing issue when reading a fragged file via directIO.• Prevents a panic when more metadata file space is needed and the disk write to allocate it fails.• Removes a restriction where dynamic VMEbus device drivers could only probe one controller per driver. Multiple controllers per driver now configure successfully.• Fixes kernel memory faults caused by ufs_sync_int accessing an inactivated or de-allocated vnode. This change also corrects a problem with negative block number detection in ufs_stratgy.• Provides support for the pseudo device /dev/poll to the kernel which allows for very efficient polling of a large number of file descriptors.• Corrects a problem that caused the RFC 2001 Fast Retransmit Algorithm within the kernel to work incorrectly.• Fixes three problems with the ee driver for DE60x Ethernet cards. These problems affect all Tru64 systems containing DE60x network interfaces.<ul style="list-style-type: none">– Fixes a race condition that can cause a panic when a transmit timeout occurs.– Improves error checking when allocating buffers.– Fixes DMA resource allocation to prevent a panic when a machine runs low on DMA resources.
---------------------------	--

Table 2–2: Summary of Base Operating System Patches (cont.)

Patch 846.00 continued	<ul style="list-style-type: none">• Adds a new ee subsystem attribute <code>link_check_interval</code> that allows the link state polling interval to be tuned for faster failover times when using DE60x interfaces for Link Aggregation.• Fixes a DMAPI problem where showfile can show that DMAPI regions exist when they do not.• Supports a cluster-specific patch which fixes a race between cluster mounts and file system lookups, and fixes a situation in which file system failover deadlocks.• Fixes an NFS readahead performance problem where performance is degraded when reading past two gigabytes in a file.• Fixes a problem where advscan incorrectly processes concatenated options (e.g., <code>-ar</code> vs. <code>-a -r</code>). For instance, if <code>-ar</code> is specified, the <code>(-r)</code> option will not be processed.• Prevents a lock hierarchy violation from occurring when AdvFS tries to extent a file on a system that is out of memory.• Addresses the problem of applications hanging with outstanding I/O during high volume I/O in a cluster environment.• Provides I/O resiliency improvements and prevents some AdvFS domain panics due to inadequate error handling between the HSG80 and the Tru64 disk driver.• Updates the EMX driver to v2.03 and fixes a problem which could cause an EMX driver panic during adapter resets.• Re-enables <code>mountd</code> to support exports file with multiline entry using leading spaces as continued line indicator. The problem was introduced with a patch that increases support of NFS file mounting from 254 to 1024 entries.• Corrects a problem with arp messages not being sent on interface static routes.• Provides the PCI indictment for storage component location to diagnose a PCI adapter failure.• Resolves a deadlock problem as well as a potential problem with incorrect or inconsistent cluster devts that could occur in a cluster when removing or replacing a device.• Corrects a problem where moving the power supply from one slot to another can cause a panic.• Corrects a possible panic when auditing <code>execve</code> with <code>exec_argp/exec_envp</code> enabled.• Allows the device special file instance numbers to be reduced to their lowest possible value and avoid runaway device names.• Fixes a problem where, under certain conditions, invalidating a portion of a very large file can make the filesystem appear to be hung. Any program trying to access the filesystem, <code>ls</code> for example, will hang until the file is invalidated. This will only happen when <code>rt_preempt_opt=1</code>.• Allows multiple applications utilizing RAID Services to send maintenance commands without interfering with each other.• Prevents different threads on multiple RADs from creating multiple references to the same level 3 page table.• Fixes <code>new_wire_method</code> (light weight wiring) issues known as the Oracle connect problem or Oracle performance problem.• Corrects an ARMTech kernel <code>malloc</code> invalid size panic.
---------------------------	---

Table 2–2: Summary of Base Operating System Patches (cont.)

Patch 846.00 continued	<ul style="list-style-type: none">• Addresses an issue with regards to creation of a crashdump on certain systems that use granularity hints regions.• Increases the number of filesystems that can be mounted from 256 to 1024. It also fixes audit to generate exportfs_create audit records correctly.• Prevents a situation in which the disk drives on a DS25 overheat if the system door is removed for too long.• A kernel memory fault panic could occur in irefresh while walking the mounted vnode list.• Fixes a problem resulting in a system panic for applications that directly call <code>nxm_get_bindings</code>.• Allows users other than root to now mount CDROM media on directories that they own.• Fixes a bug that prevented AdvFS from working correctly with LSM volumes between 1Tb and 2Tb. Also, <code>mkfdmn</code> and <code>addvol</code> will now issue a warning if an attempt is made to use an LSM volume greater than 2Tb in size.• Corrects a potential loss of data in files opened for directIO when writing in increments smaller than 8 K.• A potential security vulnerability has been discovered where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file access. We have corrected this potential vulnerability.• Fixes a problem in network Link Aggregation (LAG) where the MAC address of a LAG interface would change if the link from which it had derived its MAC address went down.• Fixes a kernel panic with a "get_xm_page_range_info:kernel memory fault" message.• Addresses several issues with page faults and stack object growth.• Corrects a problem where <code>df</code> was showing negative values for large NFS file systems.• Corrects a problem where multi-threaded processes may hang in timed condition waits (<code>pthread_cond_timedwait()</code>) when running realtime system contention scope threads.• Enables a larger maximum 1,073,741,824 for the <code>inode_hash_size</code> attribute in the UFS subsystem.• Fixes a problem that was causing the <code>tcp_rad_fasttimo</code> timer to consume excessive amounts of CPU time.• When mounting an AdvFS filesystem with the <code>-o dual</code> option in a cluster, the domain ID was not always being updated. If the mount <code>-o dual</code> happened on a node other than the node that was serving the original domain, AdvFS did not detect that the domain ID was already active and failed to update the ID for the new domain. The fix is to always create a new domain ID when the <code>-o dual</code> option is used.• Corrects a problem introduced in a prior patch which can result in a system panic when outputting through the packet filter.• This problem is seen when previously failed paths to a device (for example, <code>hsz80</code>) are restored. It can cause panics of cluster members.• Fixes a problem where AdvFS domains from a pre-V5.0 system that are mounted on a V5.1A Patch Kit 3 system would give "corrupted directory entry size" error messages when some files were accessed.
---------------------------	--



Summary of TruCluster Software Patches

This chapter summarizes the TruCluster software patches included in Patch Kit-0003.

Table 3–1 lists patches that have been updated.

Table 3–2 provides a summary of patches.

Table 3–1: Updated TruCluster Software Patches

Patch IDs	Change Summary
Patches 150.00, 195.00	New
Patches 11.00, 62.00, 97.00, 145.00, 146.00	Superseded by Patch 148.00
Patches 41.00, 80.00, 173.00	Superseded by Patch 175.00
Patches 39.00, 131.00, 178.00, 179.00	Superseded by Patch 181.00
Patches 37.00, 82.00, 132.00, 134.00, 182.00, 183.00	Superseded by Patch 185.00
Patches 70.00, 186.00	Superseded by Patch 188.00
Patches 44.00, 46.00, 189.00, 190.00, 191.00	Superseded by Patch 193.00
Patch 50.00	Superseded by Patch 200.00
Patches 12.00, 13.00, 14.00, 15.00, 16.00, 17.00, 18.00, 19.00, 20.00, 21.00, 22.00, 23.00, 25.00, 76.00, 92.00, 98.00, 99.00, 100.00, 101.00, 102.00, 103.00, 104.00, 105.00, 106.00, 107.00, 108.00, 109.00, 110.00, 111.00, 112.00, 113.00, 114.00, 116.00, 140.00, 142.00, 64.00, 86.00, 117.00, 119.00, 43.00, 151.00, 152.00, 153.00, 154.00, 155.00, 156.00, 157.00, 158.00, 159.00, 160.00, 161.00, 162.00, 163.00, 164.00, 165.00, 166.00, 167.00, 168.00, 169.00, 170.00, 172.00, 30.00, 31.00, 32.00, 33.00, 35.00, 78.00, 90.00, 122.00, 123.00, 124.00, 125.00, 126.00, 127.00, 129.00, 144.00, 196.00, 198.00	Superseded by Patch 202.00

Table 3–2: Summary of TruCluster Patches

Patch IDs	Abstract
Patch 9.00 TCR520-019	<p>Patch: Fixes networking issues within cluster environment State: Supersedes patches TCR520-008 (6.00), TCR520-037 (7.00) This patch fixes the following problems:</p> <ul style="list-style-type: none">• Multiple networking issues within a cluster environment:<ul style="list-style-type: none">– Cluster member loses connectivity with clients on remote subnets.– aliasd not handling multiple virtual aliases in a subnet and/or IP aliases.– Allows cluster members to route for an alias without joining it.– aliasd writing illegal configurations into gated.conf.memembrX.– Default route not being restored after network connectivity issues.– Fixes a race condition between aliasd and gated.– Fixes a problem with a hang caused by an incorrect /etc/hosts entry.• Fixes aliasd_niff to allow EVM restart.• Provides enablers for the Compaq Database Utility.
Patch 27.00 TCR520-028	<p>Patch: Fix for clusterwide wall messages not being received State: Existing This patch allows the cluster wall daemon to restart following an EVM daemon failure.</p>
Patch 52.00 TCR520DX-001	<p>Patch: Fixes smsd/caad performance problems State: Existing This patch provides enablers for the Compaq Database Utility.</p>
Patch 68.00 TCR520-045	<p>Patch: Fix for confusing panics on SMP systems State: Existing This patch fixes a problem where node reboots during a clusterwide shutdown would result in difficult to diagnose system panics.</p>
Patch 88.00 TCR520-076	<p>Patch: Fix for cluster hang during boot State: Supersedes patch TCR520-027 (29.00) This patch addresses a situation where the second node in a cluster hangs upon boot while setting the current time and date with ntpdate.</p>

Table 3–2: Summary of TruCluster Patches (cont.)

Patch 95.00 TCR520-071	<p>Patch: Fix for CAA problems</p> <p>State: Supersedes patches TCR520-029 (1.00), TCR520-035 (2.00), TCR520-022 (3.00), TCR520-032 (5.00), TCR520-054 (53.00), TCR520-047 (54.00), TCR520-048 (55.00), TCR520-051 (56.00), TCR520-056 (57.00), TCR520-046 (58.00), TCR520-052 (60.00), TCR520-049 (66.00), TCR520-065 (71.00), TCR520-060 (72.00), TCR520-063 (74.00), TCR520-072 (84.00), TCR520-102 (93.00)</p> <p>This patch corrects the following:</p> <ul style="list-style-type: none">• Increases parallelism in CAA event handling.• CAA cannot start or stop resources. The resource moves to the unknown state. Also, a core file is left behind by the action of starting and stopping resources. The problem will occur after the first resource is started.• Enables the Compaq Database Utility.• The datastore may get corrupted due to improper datastore locking. This may occur when multiple CAA CLI commands are run in the background.• The <code>caa_profile</code> command may complain of failure to create and log EVM events.• The <code>caa_profile -create</code> command inserts extra attributes such as REBALANCE into the profile when a user uses it to create an application profile. This will cause CAA GUI to fail to validate the profile.• The <code>caa_stat</code> command can crash, leaving a core file, when it receives a SIGPIPE signal. The problem has been known to occur when <code>caa_stat</code> output is piped to a command such as <code>head</code>.• When long resource or attribute names are used the space will not be reclaimed correctly when the resource is unregistered.• Fixed a caad memory leak caused by <code>caa_stat -f</code>.• CAA fails to close a TDF after processing a corresponding resource profile. Over time this will lead to reaching the process limit for open file descriptors and will prevent CAA from functioning properly.• The <code>clu_mibs</code> agent has been changed to retry the connection with the Event Manager daemon (<code>evmd</code>) indefinitely until it succeeds.• the <code>clu_mibs</code> agent's start and stop control has been moved from <code>/sbin/init.d/clu_max</code> script to <code>/sbin/init.d/snmpd</code> script.• Resolves erroneous behavior of resources with dependencies upon other resources (required resources). This solves several problems with starting, stopping, and relocating a resource with dependencies when the resource's start or stop scripts fail, or when relocating during a shutdown.• Migrates the old datastore to the new datastore during the rolling upgrade and corrects the problem where no resource information was preserved.• Resolves the issue with the default CAA system services (<code>dhcp</code> named <code>cluster_lockd</code> autofs) not running after the installation of the patch kit. In addition to the default CAA system services, any previously registered resource would be lost.• Prevents member hangs during boot in unusual circumstances that cause the CAA daemon to crash or exit during initialization.• Fixes three CAA problems triggered by heavy CAA activity conditions.
---------------------------	---

Table 3–2: Summary of TruCluster Patches (cont.)

Patch 121.00 TCR520-114	Patch: Using a cluster as a RIS server causes panic State: New This patch corrects the following: <ul style="list-style-type: none">• A panic caused by a known problem, using a cluster as a RIS server.• A fix to RIS/DMS serving in a TruCluster.
Patch 136.00 TCR520-085	Patch: Enhancement for clu_autofs shutdown script State: Existing This patch makes the /sbin/init.d/clu_autofs script more robust.
Patch 138.00 TCR520-121	Patch: Provides enhanced clu_upgrade switch State: Supersedes patch TCR520-009 (48.00) This patch corrects the following: <ul style="list-style-type: none">• Provides a warning to users who have installed a patch kit that includes a patch which requires a version switch. The warning informs the user that the installed patches include a version switch which cannot be removed using the normal patch removal procedure. The warning allows the user to continue with the switch stage or exit clu_upgrade.• Provides additional user information after the user has decided to perform a patch rolling upgrade and has entered the pathname to a patch kit which contains one or more patches requiring a version switch. The additional user information identifies the patches containing the version switch and provides references to the appropriate user documentation.• Addresses a problem seen during the setup stage of a rolling upgrade during tag file creation. The fix is to change a variable to only look at 500 files at a time while making tag files, instead of the current 700.
Patch 148.00 TCR520-134	Patch: Fixes cluster hang during Memory Channel initialization State: Supersedes patches TCR520-013 (11.00), TCR520-055 (62.00), TCR520-106 (97.00), TCR520-132 (145.00), TCR520-152 (146.00) This patch corrects the following: <ul style="list-style-type: none">• Fixes a situation in which one or several cluster members would panic if a Memory Channel cable was removed or faulty.• Fixes the following problems with Memory Channel in a cluster environment:<ul style="list-style-type: none">– A problem with the Memory Channel power off in LAN interconnect cluster which causes a clusterwide panic.– A user is now allow to kill a LAN interconnect cluster via Memory Channel.– Supports Memory Channel usage in a LAN cluster.• Corrects a problem when the master failover node goes off line during a failover and fails over due to parity errors increasing beyond the limit.• Addresses a problem in which a bad Memory Channel cable causes a cluster member to panic with a panic string of "rm_eh_init" or "rm_eh_init_praill".• Contains changes that will make Memory Channel failovers work better, and will also handle bad optical cables.• Fixes a problem in which a node booting into a cluster hangs during Memory Channel initialization.

Table 3–2: Summary of TruCluster Patches (cont.)

Patch 150.00 TCR520-142	Patch: Eliminates spurious duplicate error message State: New This patch eliminates a spurious duplicate error message.
Patch 175.00 TCR520-128	Patch: Resolves issues with version switched patches State: Supersedes patches TCR520-024 (41.00), TCR520-057 (80.00), TCR520-154 (173.00) This patch corrects the following: <ul style="list-style-type: none">• Fixes a cluster installation problem of having an LSM disk and a disk media with the same name. Normally, the install script would not let you install because it was looking at the disk name, not the disk media name. This has been fixed.• Disks over 10 GB are unable to be used as member or quorum disks. This fix allows the user to use them as such.• Resolves issues with version-switched patches and cluster installation. Previously, the user could run with old functionality if they had not run versw; now dupatch automatically runs it for them.• Automatically enables ip filtering for the cluster interconnect on cluster installation and member addition; allows installation on unlabeled disks; and allows the cluster installation to detect layered product kits in /var as well as /usr/var.
Patch 181.00 TCR520-115	Patch: Fixes problems in the DLM subsystem State: Supersedes patches TCR520-034 (39.00), TCR520-074 (131.00), TCR520-123 (178.00), TCR520-122 (179.00) This patch corrects the following: <ul style="list-style-type: none">• Fixes a panic in DLM when another node in the cluster is halted.• Fixes a panic in the DLM deadlock detection code.• Fixes a problem where a process using the Distributed Lock Manager can take up to ten minutes to exit.• Fixes several DLM related crashes and performance issues.• Corrects a cluster member panic.• DLM was not always returning the resource block information for the sublock even if the sublock was held.

Table 3–2: Summary of TruCluster Patches (cont.)

Patch 185.00 TCR520-125	<p>Patch: Resolves an RDG panic in the RdgShutdown routine</p> <p>State: Supersedes patches TCR520-015 (37.00), TCR520-058 (82.00), TCR520-087 (132.00), TCR520-105 (134.00), TCR520-141 (182.00), TCR520-150 (183.00)</p> <p>This patch corrects the following:</p> <ul style="list-style-type: none">• Enables the Compaq Database Utility.• Changes RDG wiring behavior to match the VM fix to wiring GH chunks.• The RDG fix closes a timing window that can cause Oracle 9i to hang when a remote node in the cluster goes down.• Fixes a possible panic on process termination and a panic involving multiple Memory Channel adapters.• Makes the RDGinit daemon program safe to execute multiple times on all cluster interconnect types.• Resolves a problem resulting in an incorrect error status being returned from RDGinit.• Fixes a Reliable DataGram (RDG) problem that can result in user processes hanging in an uninterruptable state.• Resolves an RDG panic in the RdgShutdown routine.
Patch 188.00 TCR520-138	<p>Patch: Fixes cluster kernel problem that causes a hang</p> <p>State: Supersedes patches TCR520-042 (70.00), TCR520-133 (186.00)</p> <p>This patch corrects the following:</p> <ul style="list-style-type: none">• Fixes a panic in the kernel group services when another node is booted into the cluster.• Fixes a problem in the cluster kernel that causes the cluster to hang when a member is rebooted into the cluster.• Fixes a problem in the cluster kernel that causes one or more members to panic during a cluster shutdown.
Patch 193.00 TCR520-146	<p>Patch: Fix for ICS_UNABLE_TO_MAKE_PROGRESS panic</p> <p>State: Supersedes patches TCR520-021 (44.00), TCR520-023 (46.00), TCR520-139 (189.00), TCR520-145 (190.00), TCR520-127 (191.00)</p> <p>This patch corrects the following:</p> <ul style="list-style-type: none">• Fixes a situation where ICS is unable to make progress because heartbeat checking is blocked or the input thread is stalled. The symptom is a panic of a cluster member with the panic string ICS_UNABLE_TO_MAKE_PROGRESS: HEARTBEAT CHECKING BLOCKED/INPUT THREAD STALLED.• Fixes the problem of a cluster member failing to rejoin the cluster after Memory Channel failover.• Addresses a panic that occurs when higher priority threads running on a cluster member block the internode communication service Memory Channel transport (ics_ll_mct) subsystem's input thread from execution.• Fixes numerous panics and hangs with the way a cluster communicates with its nodes. It also fixes hangs and panics during boot.• Fixes a panic with the string "rcnx_status: different node."• Fixes a boot hang on "ics_mct: Node arrival waiting for out of line node down cleanup to complete".

Table 3–2: Summary of TruCluster Patches (cont.)

Patch 195.00 TCR520-143	Patch: Memory Channel API problem causes system hang State: New This patch fixes a problem in the Memory Channel API that can cause a system to hang.
Patch 200.00 TCR520-137	Patch: Fix for ICS_BROADCAST_SETUP panic State: Supersedes patches TCR520-025 (50.00) This patch corrects the following: <ul style="list-style-type: none">• Fixes a situation where a cluster shutdown under load on a cluster using a LAN interconnect takes a very long time.• On boot, "duplicate incoming connections" will not cause a panic. Provides a complete and better error message in event of a misconfigured ICS/TCP adapter.
Patch 202.00 TCR520-167	Patch: Fixes several Device Request Dispatcher problems State: Supersedes patches (12.00), TCR520-011 (13.00), TCR520-005 (14.00), TCR520-002 (15.00), TCR520-004 (16.00), TCR520-039 (17.00), TCR520-014 (18.00), TCR520-016 (19.00), TCR520-018 (20.00), TCR520-010 (21.00), TCR520-012 (22.00), TCR520-026 (23.00), TCR520-001 (25.00), TCR520-068 (76.00), TCR520-100 (92.00), TCR520-090 (98.00), TCR520-091 (99.00), TCR520-104 (100.00), TCR520-080 (101.00), TCR520-083 (102.00), TCR520-089 (103.00), TCR520-095 (104.00), TCR520-099 (105.00), TCR520-078 (106.00), TCR520-101 (107.00), TCR520-081 (108.00), TCR520-082 (109.00), TCR520-070 (110.00), TCR520-092 (111.00), TCR520-059 (112.00), TCR520-062 (113.00), TCR520-093 (114.00), TCR520-084 (116.00), TCR520-136 (140.00), TCR520-116 (142.00), TCR520-053 (64.00), TCR520-067 (86.00), TCR520-044 (117.00), TCR520-077 (119.00), TCR520-003 (43.00), TCR520-159 (151.00), TCR520-124 (152.00), TCR520-126 (153.00), TCR520-109 (154.00), TCR520-119 (155.00), TCR520-147 (156.00), TCR520-118 (157.00), TCR520-110 (158.00), TCR520-111 (159.00), TCR520-120 (160.00), TCR520-130 (161.00), TCR520-112 (162.00), TCR520-113 (163.00), TCR520-107 (164.00), TCR520-131 (165.00), TCR520-151 (166.00), TCR520-108 (167.00), TCR520-144 (168.00), TCR520-157 (169.00), TCR520-148 (170.00), TCR520-129 (172.00), TCR520-033 (30.00), TCR520-017 (31.00), TCR520-006 (32.00), TCR520-007 (33.00), TCR520-020 (35.00), TCR520-064 (78.00), TCR520-075 (90.00), TCR520-079 (122.00), TCR520-094 (123.00), TCR520-096 (124.00), TCR520-097 (125.00), TCR520-088 (126.00), TCR520-098 (127.00), TCR520-103 (129.00), TCR520-155 (144.00), TCR520-149 (196.00), TCR520-162 (198.00) This patch corrects the following: <ul style="list-style-type: none">• Provides the I/O barrier code that prevents HSG80 controller crashes (firmware issue).• Fixes a situation in which a rebooting cluster member would panic shortly after rejoining the cluster if another cluster member was doing remote disk I/O to the rebooting member when it was rebooted.• Allows high density tape drives to use the high-density compression setting in a cluster environment.• Fixes a kernel memory fault panic that can occur within a cluster member during failover while using shared served devices.• Fixes the problem of cluster-wide hang because of DRD node failover is stuck and unable to bid a new server for served device.• Adds DRD Barrier retries to fixes for HSx firmware problems.

Table 3–2: Summary of TruCluster Patches (cont.)

Patch 202.00 continued	<ul style="list-style-type: none">• Fixes a problem where CAA applications using tape/changers as required resources will not come on line (as seen by caa_stat).• Fixes a problem where the tape changer is only accessible from the member that is the DRD server for the changer.• Fixes a problem where an open request to a disk in a cluster fails with an illegal errno (≥ 1024).• Fixes a problem where an open to a tape drive in a cluster would take 6 minutes (instead of 2) to fail if there were no tape in the drive.• Solves a problem in which a cluster would hang the next time a node was rebooted after a tape device was deleted from the cluster.• Fixes a domain panic in a cluster when a file system is mounted on a disk accessed remotely over the cluster interconnect.• Fixes the race condition problem when multiple unbarrierable disks failed at the same time.• Fixes a kernel memory fault in drd_open.• Prevents an infinite loop in drd_open().• Fixes several Device Request Dispatcher problems.• Removes a rolling upgrade issue with CDROM and FLOPPY device handling.• Addresses a problem in which a cluster or a device can get blocked I/O, or a cluster node may panic after a device has been deleted.
---------------------------	---

Table 3–2: Summary of TruCluster Patches (cont.)

Patch 202.00 continued	<ul style="list-style-type: none">• Makes AdvFS fileset quota enforcement work properly on a cluster.• Corrects a "cfsdb_assert" panic which can occur following the failure of a cluster node.• Corrects a problem which can cause cluster members to hang waiting for the update daemon to flush /var/adm/pacct.• Prevents a potential hang that can occur on a CFS failover.• Allows POSIX semaphores/msg queues to operate properly on a CFS client.• Addresses a potential file inconsistency problem which could cause erroneous data to be returned when reading a file at a CFS client node. There is also a small possibility that this problem could result in a CFS panic ("AssertFailed: bp->b_dev").• Addresses two potential CFS panics that might occur for a DMAPI/HSM managed filesystem. The first panic problem string is: Assert Failed: (t)->cntk_mode <= 2" The second panic problem string is: Assert Failed: get_recursion_count(current_threa&CFS_CMI_TO_REC_LOCK(mi)) == 1• Addresses a possible panic which could occur if multiple CFS client nodes leave the cluster while a CFS relocate or unmount is occurring.• Addresses a possible KMF panic when executing the command cfsmgr -a DEVICES on a filesystem with LSM volumes.• Corrects a CFS problem that could cause a panic with the panic string of "CFS_INFS full".• Addresses a potential CFS panic that might occur when a file being opened in direct I/O mode, while at the same time the file is being truncated by a separate process.• Provides enabler support for the Enterprise Volume Manager product.• Fixes a memory a leak in cfscall_ioctl().• Provides freezefs support.• Addresses a data inconsistency that can occur when a CFS client reads a file that was recently written to and whose underlying AdvFS extent map contains more than 100 extents.• Fixes a panic that would occur during the mount of a clustered file system on top of a nonclustered file system.• Prevents a Kernel Memory Fault panic during unmount in a cluster or during a planned relocation.• Fixes support for mounting other filesets from the cluster_root domain in a cluster.• Fixes the assertion failure ERROR != ECFS_TRYAGAIN.
---------------------------	--

Table 3–2: Summary of TruCluster Patches (cont.)

Patch 202.00 continued	<ul style="list-style-type: none">• Fixes a race condition during cluster mount which results in a transient ENODEV seen by a name space lookup.• Fixes a possible panic on boot if mount request is received from another node too early in the boot process.• Fixes a PANIC: CFS_ADD_MOUNT() - DATABASE ENTRY PRESENT panic when a node re-joins the cluster.• Fixes two race conditions in Cluster Mount support:<ul style="list-style-type: none">– One results in a transient mount failure.– The second might result in a kernel memory fault panic during mount.• Fixes a cluster problem with hung unmounts (possibly seen as hung node shutdowns).• Addresses a potential UBC panic which could occur when accessing CFS filesystems.• Fixes a possible Kernel Memory Fault panic on racing mount update/unmount/remount operations for the same mount point.• Fixes a possible race between node shutdown and unmount.• Fixes a possible Kernel Memory Fault panic on the mount update on a Memory File System (MFS) and other possible panics when bad arguments are passed to the mount library interface.• Prevents a panic “Assert failed: vp->v_numoutput > 0” or a system hang when a filesystem becomes full and direct async I/O via CFS is used. A vnode will exist that has v_numoutput with a greater than 0 value and the thread is hung in vflushbuf_aged().• Fixes a possible Kernel Memory Fault in function ckidtokgs.• Fixes a potential CFS deadlock.• Correct a cfsmgr error "Not enough space" when attempting to relocate a file system with a large amount of disks.• Addresses possible CFS client node file read failures which could occur if on a previous failure to perform a failover mount on the client node the domain storage devices were closed.• Fixes support for mounting other filesets from a cluster node’s boot partition domain.• Addresses a cluster problem that can arise in the case where a cluster is serving as an NFS server. The problem can result in stale data being cached at the nodes which are servicing NFS requests.• Addresses a CFS panic that might occur for a DMAPI/HSM managed fs: (panic): cfstok_hold_tok(): held token table overflow
---------------------------	---
