Tru64 UNIX 5.1 and TruCluster Server 5.1

Patch Summary and Release Notes for Patch Kit-0005

April 2002

This manual describes the release notes and contents of Patch Kit-0005. 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*.

Compaq Computer Corporation
Houston, Texas

Contents

About This Manual

This manual contains information specific to Patch Kit-0005 for the Tru64™ UNIX 5.1 operating system and TruCluster Server Software™ 5.1 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

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

You can send your comments in the following ways:

- Fax: 603-884-0120 Attn: UBPG Publications, 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 Compaq technical support office. Information provided with the software media explains how to send problem reports to Compaq.

1

Release Notes

This chapter provides important information that you need in order to work with the Tru64 UNIX 5.1 and TruCluster 5.1 Patch Kit-0005.

1.1 Patch Process Resources

Compaq provides Web sites to help you with the patching process:

- To obtain the lastest patch kit for your operating system and cluster:

  http://ftp1.support.compaq.com/public/unix/

- To view or print the lastest 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/index.html

- To visit Compaq's 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. Compaq recommends 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 ~88 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 ~88 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 ~129 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

Note

A rolling upgrade has specific disk space requirements. Be sure to check your disk space before starting a rolling upgrade. Make sure that your system contains the required space in all file systems before you

begin the setup stage of the roll. If any file system fails to meet the minimum space requirements, the program will fail and generate an error message similar to the following:

```
***Error***
The tar commands used to create tagged files in the '/' file system have
reported the following errors and warnings:
NOTE: CFS: File system full: /

        tar: sbin/lsm.d/raid5/volsd : No space left on device
        tar: sbin/lsm.d/raid5/volume : No space left on device
NOTE: CFS: File system full: /

.NOTE: CFS: File system full: /
```

If you receive this message, run the `clu_upgrade -undo setup` command, free up or add the required amount of space on the affected file systems, and then rerun the `clu_upgrade setup` command.

Rolling upgrade disk space requirements are described in Section 7.4.1 of the *TruCluster Server Software Installation* manual.

- Temporary Storage Space

  A total of ~250 MB of storage space is required to untar this patch kit. Compaq recommends 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 ~50 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 ~52 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 ~1194 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 Baselevel 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 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 patch kit 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
```

1.5 Release Note for KZPCC

Under heavy I/O conditions, an `open()` call to the KZPCC driver can return an I/O error. If this occurs, add the following stanza to your `sysconfigtab` file:

```
I2O:
    Max_Job_Pool_Size=1024
```

In addition, a KZPCC system can hang if you do a physical I/O greater than 4 MB. This is more likely to occur doing I/O to a raw disk with large block size transfers, but can also occur on block devices.

1.6 Release Note for Tru64 UNIX Patch 921.01

1.6.1 Removal of the directio Cloning Patch

This patch provides a script that will allow a user to remove the `directio` cloning patch after the version switch has been thrown by running `clu_upgrade -switch`. 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 `dupatch`.

The `/usr/sbin/clone_versw_undo` script must be run by root in multiuser mode after the `directio` cloning patch has been completely rolled in and before another rolling upgrade has begun. A system or cluster shut down will be required to remove the `directio` cloning patch.

Note

Since 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 responsiblity 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 which attempt to reboot before the entire cluster is shutdown can experience panics or hangs.

See the *Patch Kit Installation Instructions* for further information.

1.6.2 AdvFS and Direct I/O

In laboratory testing, Compaq has observed that under certain circumstances, a possibility exists that inconsistent data may be written to disk on some Tru64 UNIX V5.0A and V5.1 systems running AdvFS and direct I/O.

Compaq became aware of this possibility only during laboratory testing. To our knowledge, no customer has experienced this problem. Compaq is alerting customers to this potential problem as a precautionary measure.

The conditions under which this potential problem may occur are as follows:

- An application writes to a file using AdvFS direct I/O and the file had previously been opened for normal I/O (which by default is cached).
- Some but not all of the pages are still resident in Unified Buffer Cache (UBC) memory.

Invalid data could occur when a single direct I/O write spans multiple AdvFS pages, and some, but not all, of the pages are still in the UBC. If the file has been opened only for direct I/O and remains open for direct I/O, the problem does not exist.

Applications that use direct I/O, such as Oracle, could be affected.

Configurations Affected

The potential problem may affect the following systems:

- Tru64 UNIX V5.0A clustered and nonclustered systems
- Tru64 UNIX V5.1 nonclustered systems only

Only V5.0A and V5.1 systems running an application that uses direct I/O could experience this potential problem. Any application using direct I/O must request this feature explicitly.

The following Oracle versions use direct I/O and may therefore be affected:

- Oracle 8.1.7
- Oracle 8.1.6.3
- Oracle 8.1.6.2 with patch 1527141
- Oracle 8.0.6.2 with patch 1523186
- Oracle 7.3.4.5 with patch 1523179

In addition, the AdvFS file system that is used for any of the following Oracle files:

- Control file
- Data file
- Log file

An Oracle environment meeting the above criteria could experience this potential problem.

Oracle running on raw partitions exclusively or running LSM on raw partitions exclusively are not affected.

Some customers write their own applications that use direct I/O. These customers should be aware of the detailed circumstances under which this problem could occur. The problem could occur as follows:

- The write spans multiple AdvFS 8K pages.
- The last page to be written is in the UBC.
- One or more of the preceding pages are not in the UBC.
- The write to the last page is less than a full page size (8K).

Under these circumstances, the data written at the start of the total write is the original data, offset by the amount of data written to the last page.

Tru64 UNIX versions V4.* and V5.0 are NOT affected.

The potential problem is fixed in future Tru64 UNIX versions and in V5.0 Patch Kit 3 and V5.1 Patch Kit 3.

Problem

If Oracle customers are running one of the affected Oracle configurations, Oracle may have already detected an inconsistency in the database and reported errors similar to the following in the alert log and trace file:

```
ORA-01578: ORACLE data block corrupted (file # 1, block # 100)
ORA-01119: data file 1: '/scratch/820/qa/dbs/t_db1.f'


ORA-00368: checksum error in redo block
ORA-00354: Log corruption near block #231
```

Oracle customers that have run the dbverify (dbv) utility may have encountered an error message similar to the following:

```
***
  Corrupt block relative dba: 0x0040900b (file 0, block 36875)
  Bad header found during dbv:
  Data in bad block -
   type: 27 format: 2 rdba: 0x0040900d
   last change scn: 0x0000.0001349a seq: 0x2 flg: 0x04
   consistency value in tail: 0x349a1b02
   check value in block header: 0xa377, computed block checksum: 0x0
   spare1: 0x0, spare2: 0x0, spare3: 0x0
  ***
```

1.6.3 Technical Update for KZPCC products

This patch provides support for KZPCC products.

For more information see Tru64 UNIX technical updates provided at the following URL:

http://www.tru64unix.compaq.com/faqs/publications/patch/

Select the option for Operating System Technical Updates and choose the following document:

Tru64 UNIX Version 5.1 Technical Update

This technical update will also contain information for valid upgrade paths to Tru64 UNIX Version 5.1 from the Version 4.0x releases that currently support I2O.

1.6.4 Problem with Multi-user Mode Application

Warning

When applying this patch in multi-user mode, an inconsistency problem results between the updated `/shlib/libpthread.so` and the existing kernel. The problem manifests itself when you install the patch in multi-user mode and you elect to reboot at a later time. The scheduled reboot will not occur. This problem can be avoided by installing Patch 921.01 in single user mode, or selecting the option to reboot now (rather than scheduling later).

To correct this situation, if you have installed the patch and have not rebooted the system, execute the following commands:

1. Set `DUPATCH_SESLOG` to location of session log, by default:

   `/var/adm/patch/log/session.log`

2. Get the name of newly-built kernel:

```
# NEW_KERNEL=‘grep "The new kernel is" $DUPATCH_SESLOG | awk
’ { print $5 }’ ‘
```

3. Copy the new kernel:

   ```
   # cp <NEW_KERNEL> /vmunix
   ```

4. Reboot the system at a specified time:

   ```
   # shutdown -r <TIME_OF_REBOOT>
   ```

After rebooting with the new kernel, your system will once again be consistent.

1.6.5 New Graphics Card

This patch provides the driver support for a new graphics card. In order to obtain full support for this graphics card, you must also select Patch 777.00, which is the X server portion of the patch.

A list of supported platforms is available on the following web page:

http://www.compaq.com/alphaserver/products/options.html

If you have a system with this new graphics card, you will need to reconfigure and rebuild the kernel after installing this patch.

To do this, follow these steps:

1. Shut down the system:

   ```
   # /usr/sbin/shutdown -h now
   ```

2. Boot genvmunix to single-user mode:

   ```
   >>> boot -fi genvmunix -fl s
   ```

3. After the system boots to single-user mode, mount the file systems, run the update command, and activate the swap partition:

   ```
   # /sbin/bcheckrc
   # /sbin/update
   # /sbin/swapon -a
   ```

4. Run doconfig to create a new kernel configuration file and rebuild the kernel:

   ```
   # /usr/sbin/doconfig
   ```

   Note

   Do not specify the -c option to doconfig. If you do, doconfig will use the existing kernel configuration file which will not have the appropriate controller entry for the new graphics card.

5. Save the old /vmunix file and move the new kernel to /vmunix.

6. Shut down the system:

   ```
   # /usr/sbin/shutdown -h now
   ```

7. Boot the new kernel:

   ```
   >>> boot
   ```

If you remove this patch from your system after you have rebuilt the kernel to incorporate support for the new graphics card as described previously, you will need to rebuild the kernel again to restore generic VGA graphics support. To do this, follow the steps given previously. The doconfig running on the original, unpatched genvmunix will not recognize the new graphics card and will include generic VGA graphics support in the resulting kernel.

1.6.6 DEGPA-TA Gigabit Ethernet Device

This patch provides support for DEGPA-TA (1000BaseT) Gigabit Ethernet device. If you have a system with this new Ethernet device, you will need to reconfigure and rebuild the kernel after installing this patch.

To do this, follow these steps:

1. Shut down the system:

   # **/usr/sbin/shutdown -h now**

2. Boot genvmunix to single-user mode:

   >>> **boot -fi genvmunix -fl s**

3. After the system boots to single-user mode, mount the file systems, run the update command, and activate the swap partition:

   # **/sbin/bcheckrc**

   # **/sbin/update**

   # **/sbin/swapon -a**

4. Run doconfig to create a new kernel configuration file and rebuild the kernel:

   # **/usr/sbin/doconfig**

   Note

   Do not specify the -c option to doconfig. If you do, doconfig will use the existing kernel configuration file which will not have the appropriate controller entry for the Ethernet device..

5. Save the old /vmunix file and move the new kernel to /vmunix.

6. Shut down the system:

   # **/usr/sbin/shutdown -h now**

7. Boot the new kernel:

   >>> **boot**

If you remove this patch from your system after you have rebuilt the kernel to incorporate support for the new Ethernet card as described previously, you will need to rebuild the kernel. To do this, follow the steps given previously. The doconfig running on the original, unpatched genvmunix will not recognize the new Ethernet driver.

1.6.7 Configuring FibreChannel Systems

This patch requires that FibreChannel systems which utilize FibreChannel devices for boot and swap be properly configured as follows:

- There is a minimum of 1.25 -2 times physical memory for swap space available.
- All boot and swap devices are properly configured to use one of the four console ports.
- The console WWID number for each boot or swap device is identical to the WWID number found via the hwmgr utility using the steps outlined as follows:

  1. Identify the console port(N) and WWID number configuration information using consvar as follows:

     ```
     consvar -g N1 ; consvar -g N2
     consvar -g N3 ; consvar -g N4
     consvar -g wwid0 ; consvar -g wwid1
     consvar -g wwid2 ; consvar -g wwid3
     ```

2. Find the device name by checking `etc/fstab`, using `showfdmn` for AdvFS root domains, and `swapon -s` for swap devices for each FibreChannel boot and swap device.
3. Find the HWID using the device name obtained in step 2.

   ```
   hwmgr -view dev | grep "device name from step 2 above"
   ```

   for each FibreChannel boot and swap device.
4. Find the WWID using the device name obtained in step 3.

   ```
   hwmgr -view dev | grep "device name from step 3 above"
   ```

   for each FibreChannel boot and swap device.
5. Verify that the hwmgr WWIDs from step 4 above match the WWIDs from step 1 above for each FibreChannel boot and swap device.
6. If the WWIDs do not match in step 5 then the system needs to be shut down and reconfigured using the wwidmgr utility as described in the Wwidmgr Users Manual located in the doc directory on the Firmware CDROM until you have verified that the WWID console configuration matches the system hwmgr WWID configuration using the steps described previously.

1.7 Release Note for Tru64 UNIX Patches 324.00 and 496.00

This patch delivers version V1.0-032 of the `libots3` library. Version 2.0 (or greater) of the `libots3` library is delivered with the Compaq FORTRAN Compiler, Versions V5.3 ECO1 and V5.4, or the Developers Tool Kit (DTK) (OTABASE subset). If `libots3` V2.0 (or greater) is already installed on your system, and you install this patch, you will receive the following informational message:

```
Problem installing:


- Tru64_UNIX_V5.1 / Software Development Environment Patches:

      Patch 00496.00 - Fix for problems in Compaq C compiler


       ./usr/shlib/libots3.so:

                 is installed by:

                                 OTABASE212
               and can not be replaced by this patch.

This patch will not be installed.
```

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

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

1.8 Release Note for Tru64 UNIX Patch 888.00

This release note contains a new reference page for the `fixfdmn` utility.

NAME
fixfdmn - Checks and repairs corrupted AdvFS domains

SYNOPSIS

/sbin/advfs/fixfdmn [-mtype[,type]...] [-d directory] [-v number] [-a [-c] | -n] [-s {y | n}] [domain] [fileset]

/sbin/advfs/fixfdmn -u directory domain

OPTIONS

-a Specifies that after repairing what it can, fixfdmn will attempt to activate the domain at the end of the run. This option cannot be used with the -n option.

-c Removes any clone filesets. This option is only valid if used with the -a option.

-d directory
Specifies a directory to which the message log and undo files will be written. If the -d option is not used, the message and undo log files are put in the current working directory. The message log file is named fixfdmn.<domain>.log and the two undo files are named undo.<domain>.<#> and undoidx.<domain>.<#> where # will cause a number to be appended to the filenames to make them unique. The numbers will be rotated sequentially from 0 (zero) through 9 if multiple undo files are created for the same domain. The undo file will have the same ending number as its corresponding undo index file.

-m type[,type...]
Specifies a list of types of metadata, one or more of which can be checked and repaired. The valid types are log, sbm, sync, bmt, frag, quota and files. If you specify the fileset parameter, sync, log, sbm, and bmt are made invalid types for the -m option. If you do not specify -m, the default is to check all types.

sync
Corrects the magic number and synchronizes data across volumes (for example, volume numbers, mount ids, mount states, domain ids, and so on.)

log Resets the transaction log so it is not processed.

sbm Synchronizes the sbm to the information in the bmt.

bmt Corrects the bmt.

frag
Corrects frag file groups and free lists and ensures that all file frags reside in the frag file.

quota
Checks and corrects sizes of quota files.

files
Verifies that directory metadata is correct.

-n Specifies that fixfdmn will check the domain and not do any repairs. It will report what problems were found and how it would have fixed them.

-s {y | n}
Specifies that "yes" or "no" should be answered to prompts when run from a script.

-u directory
Restores the domain to its previous state by undoing the effects of the last run of fixfdmn, using the most recent undo files in the specified directory.

-v number
Specifies the verbose mode level which controls the messages printed to stdout.

0 = Only error messages

1 = (Default) Progress, errors and summary messages

2 = Progress messages, detailed error messages, fix information and summary messages

OPERANDS

domain

The name of a corrupted domain to repair.

fileset
The name of the fileset to repair if only one fileset in this domain exhibits errors. You may tell fixfdmn to check only that fileset and not specifically look for errors in other filesets.

DESCRIPTION

The fixfdmn utility checks and repairs corrupt AdvFS domains and filesets.

The fixfdmn utility is primarily concerned with fixing problems that have a limited scope. When a large portion of the domain is corrupted, there is very little fixfdmn can do, so it will recommend restoring data from backup or running the salvage(8) command.

The fixfdmn utility uses the on-disk metadata to determine what corruptions exist in the domain. Only metadata will be repaired, as there is currently no way to check or repair the contents of users files. Only those problems which prevent mounting the domain, or would result in a domain or system panic, will be repaired.

After major areas of metadata are checked, and if a corruption was fixed, fixfdmn will prompt the user to determine if they want to continue looking for additional corruption.

If fixfdmn detects an error in a clone fileset, the clone is marked out of sync and should not be used.

If fixfdmn cannot recover the metadata for a specific file, the file may be truncated, moved, or deleted depending on the situation. The fixfdmn utility will attempt to save as much of a file as possible.

Every page fixfdmn changes will be saved to an undo file. If the user does not like the results of running fixfdmn, the user can undo the changes by running fixfdmn again with the -u option. If the file system containing the undo files runs out of space during the fixfdmn run, the user will be prompted on how to proceed. The user will have the option to continue without the undo files, to continue adding more space to the domain containing the undo files, or to exit.

Use the -m type option when you have information from a system/domain panic or output from verify or other tools which indicate where the corruption may be. This option limits the scope of what is checked and repaired.

NOTES

The fixfdmn command will always clear the transaction log, even on a non-corrupt domain unless the -n option is specified

There must be a domain entry for this domain in /etc/fdmns. The fixfdmn command opens the block devices specified for the volumes in /etc/fdmns.

If you need to repair the root domain, you must boot from CD-ROM and create the entry for the root domain under /etc/fdmns.

RESTRICTIONS

You must be root to run fixfdmn.

The fixfdmn command requires that the domain specified will have no filesets mounted.

Although fixfdmn may report success, it does not guarantee that all corruptions have been eliminated.

If a domain is mounted and written to after being repaired by fixfdmn, using the fixfdmnutility with the -u option will likely cause corruptions.

EXIT STATUS

0 (Zero)

Success.

1 Corrupt
Unable to repair all found corruptions

2 Failure
Program or system error

FILES

/etc/fdmns
Contains AdvFS domain directories and locks.

SEE ALSO

Commands: salvage(8), umount(8), verify(8), vrestore(8)

1.9 Release Note for Tru64 UNIX Patch 777.00

This patch provides the X server support for a new graphics card. In order to obtain full support for this graphic card, you must also select Patch 921.01, which is the driver portion of the patch.

A list of supported platforms is available on the following web page:

http://www.compaq.com/alphaserver/products/options.html

1.10 Release Note for Tru64 UNIX Patch 391.00

This patch contains a solution for the following issue:

Compaq has advised owners of DS10, DS10L, ES40 AlphaServers, and XP900 AlphaStations that Compaq has determined in laboratory testing that there is a theoretical possibility that during read and write operations to the floppy disk on these systems, a single byte of data may be inaccurately read or written without notice to the user or system. The potential for this anomaly exists only if floppy disk read or write operations are attempted while there is extremely heavy traffic on these Alpha systems' internal input/output busses.

Although Compaq has observed the anomaly only in laboratory tests designed to create atypical system stresses, including almost constant use of the floppy disk drive, Compaq has informed owners of the remote possibility that the anomaly could occur so that they may take precautions to prevent it.

Compaq recommends that the solution be installed by all DS10, DS10L, ES40 AlphaServers, and XP900 AlphaStation customers.

The solution to this issue is also available as an individual, manually installed patch kit named `floppy_CSP_v51.tar.gz`, available from:

http://ftp1.support.compaq.com/public/unix/v5.1

1.11 Release Note for Tru64 UNIX Patch 921.01

1.11.1 TMPDIR Environment 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
```

1.11.2 sys_check Version

The following information is for users who have installed the sys_check version 125 web kit or higher and are currently using the 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 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 is installed and will restore these symbolic links when the patch is 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.

1.12 Reference Page Information for Tru64 UNIX Patch 921.01

This release note contains reference page information for sys_check(8).

```
sys_check(8)                                    sys_check(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 inven-
    tory 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 sup-
```

port 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 the EVM(5) reference page 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. As 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 impact	Performance	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 directory 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 etcfiles 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 asys_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 an 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.

HSV controller-specific limitations include the following:

The sys_check utility uses the SANscript utility (sssu) to collect data from an Enterprise controller. This utility is included with the Enterprise Platform Kit; verify that this utility is installed in /usr/lbin and ensure that it has execute permissions.

The sys_check utility cannot dynamically determine the SAN appliance or appliances used to manage your Enterprise storage. To do so, create the file /etc/enterprise.txt with the element name, the user name, and the password (separated by colons) of the SAN appliance as shown below; these values may contain embedded spaces. Set the permissions of this file to 600.

```
element:user:password
element 1:user 1:password
```

The sys_check utility attempts to check the NetWorker backup schedule against the /etc/fstab file. For some older versions of Networker, the nsradmin 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 services.

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 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.kshin /var/recovery/sys_check.*:

```
cd /var/cluster/members/member/recovery/sys_check/
llist="advfs.stanza
vfs.stanza "
for stf in $llist; 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 pro-

cedure.

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/lbin.

FILES

/usr/sbin/sys_check
Specifies the command path.

Note

This file may be a symbolic link.

/usr/lbin/*
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

1.13 Release Note for Tru64 UNIX Patch 921.01

This section describes changes to Patch 921.01.

Updated sh, csh, 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

sh noclobber option and >| , >>| constructs Added

A `-noclobber` option similar to that already available with `csh` 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 filename is actually a symlink, the presence of the symlink satisfies the criteria `file exists` whether or not the symlink 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 symlink whose target does not exist, `sh` returns an error rather than create the file. The >>| construct will suppress these checks and create the file.

ksh noclobber Behavior Clarified

For > with `noclobber` set, `ksh` returns an error rather than overwrite an existing file. If the file name is actually a symlink, the presence of the symlink satisfies the criteria `file exists` whether or not the symlink 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 symlink to a non-existent file, `ksh` returns an error.

csh noclobber Behavior Clarified

For > with noclobber set, `csh` returns an error rather than overwrite an existing file. If the file name is actually a symlink, the presence of the symlink satisfies the criteria `file exists` whether or not the symlink 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 filename is actually a symlink to a non-existent file, `csh` returns an error. The >>! construct will suppress these checks and create the file.

Updated mkdir System Call and Command

This kit reverts the `mkdir` system call, and thus the `mkdir` command, to its Tru64 UNIX V4.n behavior with respect to symlinks. For the unusual case where a symlink is used as the very last elment of a `mkdir` path, the `mkdir` system call nows returns an erro rrather than create the target.

If you want `mkdir` to follow the symlink, you can do this by making the last character of the `mkdir` pathname a slash. The following text describes how to get `mkdir` to follow the symlink.

If /var/tmp/foo is a symlink 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 system wide by an addition to the sysconfig options for the vfs subsystem. The new sysconfig option -follow_mkdir_symlinks defaults to 0, specifying the secure symlink behavior. Changing this option to 1, which Compaq strongly discourages, will cause mkdir to follow symlinks.

1.14 Release Note for Broken Link Problem

When performing a baseline analysis with the dupatch utility on Tru64 UNIX 5.1 systems, the baseline error log files may report that a number of files have broken hard links to the /usr/share/man/man3 directory.

The presence of these broken links will not affect your system operation, the installation of dupatch or dupatch tools, the successful installation of patches, or the rebuilding of kernels on the system. The problem will be addressed in a future version of the operating system.

You can determine if these broken links exist on your system by performing the following steps:

1. Change directories as follows:

   # **cd /usr/share/man/man3**

2. Check to see that the inodes are the same for all the files:

   # **ls -il slk*.3.gz curs_slk.3.gz**

An example of a correct hard link would look as follows. Note the same inodes.

```
14648 -rw-r--r--  17 root     system      2086 Mar  9  2000 curs_slk.3.gz
14648 -rw-r--r--  17 root     system      2086 Mar  9  2000 slk_attr_off.3.gz
14648 -rw-r--r--  17 root     system      2086 Mar  9  2000 slk_attr_on.3.gz
14648 -rw-r--r--  17 root     system      2086 Mar  9  2000 slk_attr_set.3.gz
14648 -rw-r--r--  17 root     system      2086 Mar  9  2000 slk_attroff.3.gz
14648 -rw-r--r--  17 root     system      2086 Mar  9  2000 slk_attron.3.gz
14648 -rw-r--r--  17 root     system      2086 Mar  9  2000 slk_attrset.3.gz
14648 -rw-r--r--  17 root     system      2086 Mar  9  2000 slk_clear.3.gz
14648 -rw-r--r--  17 root     system      2086 Mar  9  2000 slk_color.3.gz
14648 -rw-r--r--  17 root     system      2086 Mar  9  2000 slk_init.3.gz
14648 -rw-r--r--  17 root     system      2086 Mar  9  2000 slk_label.3.gz
14648 -rw-r--r--  17 root     system      2086 Mar  9  2000 slk_noutrefresh.3.gz
14648 -rw-r--r--  17 root     system      2086 Mar  9  2000 slk_refresh.3.gz
14648 -rw-r--r--  17 root     system      2086 Mar  9  2000 slk_restore.3.gz
14648 -rw-r--r--  17 root     system      2086 Mar  9  2000 slk_set.3.gz
14648 -rw-r--r--  17 root     system      2086 Mar  9  2000 slk_touch.3.gz
14648 -rw-r--r--  17 root     system      2086 Mar  9  2000 slk_wset.3.gz
```

An example of an incorrect hardlink would look as follows. Note the different inodes.

```
54891 -rw-r--r--   2 root     system      2086 Aug 11 17:32 curs_slk.3.gz
54891 -rw-r--r--   2 root     system      2086 Aug 11 17:32 slk_attr_off.3.gz
55583 -rw-r--r--  15 root     system      2086 Aug 11 17:32 slk_attr_on.3.gz
55583 -rw-r--r--  15 root     system      2086 Aug 11 17:32 slk_attr_set.3.gz
55583 -rw-r--r--  15 root     system      2086 Aug 11 17:32 slk_attroff.3.gz
55583 -rw-r--r--  15 root     system      2086 Aug 11 17:32 slk_attron.3.gz
55583 -rw-r--r--  15 root     system      2086 Aug 11 17:32 slk_attrset.3.gz
55583 -rw-r--r--  15 root     system      2086 Aug 11 17:32 slk_clear.3.gz
55583 -rw-r--r--  15 root     system      2086 Aug 11 17:32 slk_color.3.gz
55583 -rw-r--r--  15 root     system      2086 Aug 11 17:32 slk_init.3.gz
55583 -rw-r--r--  15 root     system      2086 Aug 11 17:32 slk_label.3.gz
55583 -rw-r--r--  15 root     system      2086 Aug 11 17:32 slk_noutrefresh.3.gz
55583 -rw-r--r--  15 root     system      2086 Aug 11 17:32 slk_refresh.3.gz
55583 -rw-r--r--  15 root     system      2086 Aug 11 17:32 slk_restore.3.gz
55583 -rw-r--r--  15 root     system      2086 Aug 11 17:32  slk_set.3.gz
55583 -rw-r--r--  15 root     system      2086 Aug 11 17:32 slk_touch.3.gz
55583 -rw-r--r--  15 root     system      2086 Aug 11 17:32 slk_wset.3.gz
```

1.15 Release Note for Potential Rolling Upgrade Problem

When patching a clustered Tru64 UNIX 5.1 system using the rolling upgrade procedure, the operation may fail if your system has been upgraded from a patched Tru64 UNIX 5.0A version.

In such cases, the lead member is successfully patched, but the patching operation fails for subsequent members. The problem occurs because the file `var/adm/patch/roll/installed_patches` contains the old `OSFPAT*505` entries, which no longer exist in `./usr/.smdb`. As a result, the rolling upgrade generates error messages such as the following when subsequent members are rolled:

```
Backing up member-specific data for member: 2
 .......
 grep: can't open ./usr/.smdb./OSFPAT00018600505.inv
 grep: can't open ./usr/.smdb./OSFPAT00019200505.inv
 grep: can't open ./usr/.smdb./OSFPAT00020500505.inv
 grep: can't open ./usr/.smdb./OSFPAT00021100505.inv
 grep: can't open ./usr/.smdb./OSFPAT00016500505.inv
 grep: can't open ./usr/.smdb./OSFPAT00020000505.inv
```

The following procedures describe how to solve the problem if you discover it during a rolling upgrade or if you have not yet begun the rolling upgrade.

Rolling Upgrade Started

Perform the following steps if you issued the `clu_upgrade` command and discovered the error during the roll of the second member (designated here as member 2):

1. Halt the failing member:

   # **halt**

2. On the lead member, undo the roll:

   # **clu_upgrade undo roll 2**

3. Remove the old `OSFPAT*505` entries from `/var/adm/patch/roll/installed_patches`. Because this is a cluster-common file, you need only do this once. The remaining members can be rolled as documented in the *Patch Kit Installation Instructions*.

   a. Change to the `/var/adm/patch/roll` directory:

      # **cd /var/adm/patch/roll**

   b. Invoke an editor such as `vi` and remove any lines that contain the string `OSFPAT*505` from the file `installed_patches`:

      # **vi ./installed_patches**

4. Boot member 2 to multiuser mode and then shut down to single-user mode:

   >>> **boot**

   # **shutdown now**

5. Roll member 2:

   # **bckeckrc**

   # **clu_upgrade roll**

6. Complete the procedure as documented in the *Patch Kit Installation Instructions*.

Rolling Upgrade Not Started

Perform the following steps if you have not started a rolling upgrade:

1. Rename the `installed_patches` file and re-create it.

```
# cd /var/adm/patch/roll/
# mv ./installed_patches ./installed_patches.V50A
# touch ./installed_patches
```

2. Complete the procedure as documented in the *Patch Kit Installation Instructions*.

For information on patching your clustered system using the rolling upgrade procedure, see the *Patch Kit Installation Instructions* and the `clu_upgrade`(8) reference page.

1.16 Release Note for Tru64 UNIX Patch 169.00

In cases where the `bttape` or `btcreate` command is used to back up and restore UFS file systems, `btextract` leaves behind a `symboltable` file in the restored file system. This file, if present, will cause `btextract` to hang the next time a bootable tape is created using `btcreate` or `bttape`. The `btextract` command hangs while trying to restore the UFS file system.

To work around this problem, ensure that the file `restoresymtab?` (where `?` refers to the cluster member ID, 0 by default) is removed. Every UFS file system that was restored using `btextract` will have this file, and this file needs to be removed on each file system before running the `bttape` or `btcreate` command the next time. For example, if `/` and `/usr` are backed up, then the file will be found at `/restoresymtable0` and `/usr/restoresymtable0`, and both instances of the file need to be removed before proceeding with `btcreate` or `bttape`.

1.17 Release Note for Tru64 UNIX Patch 270.00

This patch fixes a security vulnerability (called the Brown Orifice) in Netscape Communicator Version 4.72 by updating Netscape Communicator to Version 4.75.

To determine which version of Netscape Communicator you are running, click on the Help button in the toolbar at the top of the Navigator component window, then choose the About Communicator option from the drop down menu.

You can download the latest version of Netscape Communicator for Tru64 UNIX from the Netscape Download World Wide Web site:

`http://home.netscape.com/download/index.html`

Or, from the Compaq Tru64 UNIX World Wide Web site:

`http://www.tru64unix.compaq.com/internet/download.htm`

If you are unable to upgrade to Netscape Communicator 4.75 or later, you can avoid this security vulnerability by disabling the browser's ability to run Java by following these steps:

1. Start Netscape Communicator:

   $**`/usr/bin/X11/netscape`**

2. Click on the Edit button in the toolbar at the top of the Navigator component window.

3. Click on the Preferences... option on the drop down menu that appears when the Edit button is selected. This displays the Netscape: Preferences dialog box.

4. In the window pane on the left of the Netscape: Preferences dialog box, click on the Advanced tab. This displays the advanced Communicator preferences in the dialog box.
5. If the box next to the Enable Java preference has a check mark in it, click on the box to remove the check mark. This will disable the Java programming language. Then, click on the Okay button in the Advanced preferences dialog box. (If there is no check mark in the box, you do not need to take any action.)
6. Exit Netscape Communicator by clicking on the Exit option in the drop down menu that appears when you click on the File button on the toolbar at the top of the Navigator window.

Disabling Java ensures Netscape Communicator is not vulnerable to the Brown Orifice vulnerability. You do not have to disable JavaScript.

Note

If you use the Japanese or Chinese interfaces provided in the Worldwide Language Support software, you must update the Communicator version numbers in the `/usr/lib/X11/*/app-defaults/Netscape` file if you choose to upgrade to Netscape Communicator Version 4.75 or later.

If the version numbers in these files do not match the version of Netscape Communicator installed, it will not run in the Japanese or Chinese locales.

You can download the updated files from the Compaq Tru64 UNIX World Wide Web site:

`http://www.tru64unix.compaq.com/internet/download.htm`

1.18 Release Note for Tru64 UNIX Patch 387.00

This patch modifies the `fverify` application so that it can fix files which were erroneously installed onto the system with the date of 12/31/69. This was due to a problem in the Compact Disk File System (CDFS) code that caused any file copied onto a system from a CD created in CDFS format after 1/1/2001 to have the erroneous date. The files will be corrected automatically when `fverify` is invoked by `setld`(8) during the verification phase of the software installation.

1.19 Release Note Tru64 UNIX Patch 921.01

If you have installed sendmail from the Internet Express (IX) or Open Source Internet Solutions (OSIS) layered products (IAESMTP subset), the `mailsetup` command from this subset will cause the automated patch update mechanism to fail.

Perform the following steps to fix this problem:

1. Save the IX version of mailsetup to a temporary file:

   # **`mv /usr/sbin/mailsetup /usr/sbin/mailsetup.IX`**
2. Restore the base operating system mailsetup file:

   # **`cp -p /usr/sbin/mailsetup.preIAE5.6 /usr/sbin/mailsetup`**
3. Run the patch update.
4. After the patch update has completed, save the updated base operating system file:

   # **`cp -p /usr/sbin/mailsetup /usr/sbin/mailsetup.preIAE5.6`**

5. Then restore the IX mailsetup file:

   # **mv /usr/sbin/mailsetup.IX /usr/sbin/mailsetup**

1.20 Release Note for Tru64 UNIX Patch 900.00

There was a problem in earlier releases that caused LSM to incorrectly determine the WWID of a disk in a cluster. As a result of that, fixing this problem now means that some disks that were previously identified with an incorrect WWID now may be incorrectly rejected as a clone.

The workaround is to run `volrestore` after the update installation. This assumes that you have a current `volsave` data. If you do not have current `volsave` data then you will need to restart LSM with the prior version of `vold` (`/sbin/vold`).

This may require a cluster reboot if all volumes under LSM cannot be closed//unmounted. The older `vold` should still accept the incorrectly identified disks and you can now run `volsave`. You can then return to the current `vold`, reboot if necessary, and then run `volrestore` to correct the incorrectly identified disks.

Assuming you ran `volsave` before the upgrade then the procedure is as follows:

1. Remove all of the invalid disks from LSM control:

   # **voldisk rm dskA dskB dskC** ...

2. Run volrestore:

   # **volrestore**

3. Manually start all volumes in each of the recovered diskgroups:

   # **volume -g DG1 start V1 V2 V3** ...

   # **volume -g DG2 start V1 V2 V3** ...

This procedure is only required if the system incorrectly rejects one or more disks as clones. If you do not see this behavior then you do not need to do the `volrestore` operation.

1.21 Release Note for TruCluster Patch 152.01

This patch fixes a problem that can occur when an application does a direct I/O write (an AdvFS file was opened with the O_DIRECTIO flag) or when an application performs asynchronous Direct I/Os to files using the `aio_raw` library, and the target file resides on a fileset that has been cloned.

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 multi-user mode.

---

1. Run the /usr/sbin/clone_versw_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 multi-user 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.

1.22 Release Note for TruCluster Server Software

During the switch stage of a rolling upgrade from TruCluster Server Version 5.1 to TruCluster 5.1 Patch Kit-0005, you may see the following message:

```
Initiating version switch on cluster members
.Switch already switched
```

You can safely ignore this message. The switch stage will complete successfully.

1.23 Release Note for LP9002 Support

If you have installed LP9002 FibreChannel Adapters and have built your system for the first time you need to edit your kernel configuration file in /sys/conf and add the following line to the Static Driver Definitions section:

```
#
# Static Driver Definitions
#

config_driver emx
```

When this is done you can install V5.1 Patch Kit-0004 or Patch Kit-0005. When kernel rebuilding has been done your adapters should be recognized by the operating system and should have access to your FibreChannel devices upon system reboot.

If you are installing LP9002 Adapters and have already installed V5.1 Patch Kit-0004 or Patch Kit-0005, halt your system, boot genvmunix and rebuild your kernel using doconfig but replace your existing configuration file. Copy the kernel into place (we recommend that you keep a backup until you have booted the new kernel and things are working as expected) and reboot your system.

1.24 Release Note for Cluster Alias Routing

When you use cluster alias routing make sure that the /proc file is mounted. To verify that the file is mounted, execute the df command, as follows:

```
# df /proc
```

After you execute the df /proc command you will see this output:

```
Filesystem   512-blocks        Used   Available Capacity  Mounted on
/proc                 0           0           0   100%    /proc
```

If the file is not mounted, then edit the `/etc/fstab` file as root and add the following entry:

```
/proc                     /proc                       procfs rw 0 0
```

Then execute following command as root:

```
# mount /proc
```

2

Summary of Base Operating System Patches

This chapter summarizes the base operating system patches included in Patch Kit-0005.

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 746.00, 748.00, 770.00, 772.00, 774.00, 782.00, 787.00, 798.00, 805.00, 807.00, 818.00, 829.00, 840.00, 846.00, 848.00, 850.00, 852.00, 855.00, 861.00, 866.00, 868.00, 870.00, 876.00, 880.00, 884.00, 896.00, 898.00, 900.00, 909.00, 911.00, 913.00, 915.00, 917.00, 919.00	New
Patches 13.00, 175.00, 406.00, 408.00	Superseded by Patch 655.00
Patch 76.00	Superseded by Patch 671.00
Patches 34.00, 227.00, 228.00, 230.00	Superseded by Patch 744.00
Patches 1.00, 2.00, 3.00, 5.00, 87.00, 88.00, 90.00, 233.00, 234.00, 235.00, 236.00, 237.00, 238.00, 239.00, 240.00, 241.00, 243.00, 501.00, 502.00, 504.00, 749.00, 750.00, 751.00, 752.00, 753.00, 754.00, 755.00, 756.00, 757.00, 758.00, 759.00	Superseded by Patch 761.00
Patches 310.00, 565.00, 158.00, 762.00, 763.00, 764.00	Superseded by Patch 766.00
Patches 767.00, 768.00	Superseded by Patch 770.00
Patch 246.00, 247.00, 249.00, 119.00, 287.00, 505.00, 506.00, 507.00, 509.00, 775.00	Superseded by Patch 777.00
Patch 64.00, 256.00, 257.00, 258.00, 260.00, 521.00, 783.00	Superseded by Patch 785.00
Patch 70.00, 280.00	Superseded by Patch 796.00
Patches 288.00, 290.00, 538.00	Superseded by Patch 800.00
Patches 36.00, 122.00, 124.00, 801.00	Superseded by Patch 803.00
Patches 294.00, 808.00	Superseded by Patch 810.00
Patch 296.00	Superseded by Patch 812.00
Patches 128.00, 298.00, 541.00, 543.00, 813.00, 814.00	Superseded by Patch 816.00
Patches 859.00, 778.00, 780.00	Superseded by Patch 818.00
Patches 300.00, 548.00	Superseded by Patch 820.00
Patches 134.00, 552.00, 821.00	Superseded by Patch 823.00
Patch 554.00	Superseded by Patch 825.00
Patch 556.00	Superseded by Patch 827.00
Patches 494.00, 830.00, 831.00	Superseded by Patch 833.00
Patch 142.00	Superseded by Patch 835.00
Patch 144.00	Superseded by Patch 837.00

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

Patch 838.00	Superseded by Patch 840.00
Patches 68.00, 563.00	Superseded by Patch 842.00
Patch 314.00	Superseded by Patch 844.00
Patch 853.00	Superseded by Patch 855.00
Patch 581.00	Superseded by Patch 857.00
Patch 162.00	Superseded by Patch 863.00
Patch 864.00	Superseded by Patch 866.00
Patches 112.00, 114.00	Superseded by Patch 872.00
Patches 115.00, 117.00	Superseded by Patch 874.00
Patch 347.00	Superseded by Patch 878.00
Patches 72.00, 354.00, 355.00, 356.00, 358.00, 603.00, OSF510-240 605.00	Superseded by Patch 882.00
Patches 126.00, 610.00, 612.00	Superseded by Patch 886.00
Patch 614.00	Superseded by Patch 888.00
Patches 379.00. 381.00	Superseded by Patch 890.00
Patches 145.00, 146.00, 148.00, 370.00, 624.00	Superseded by Patch 892.00
Patches 160.00. 393.00	Superseded by Patch 902.00
Patches 395.00, 638.00	Superseded by Patch 904.00

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

Patches 544.00, 546.00	Superseded by Patch 917.00
Patches 132.00, 330.00, 250.00, 252.00, 59.00, 156.00, 53.00, 60.00, 62.00, 151.00, 152.00, 154.00, 11.00, 22.00, 23.00, 24.00, 25.00, 26.00, 27.00, 28.00, 29.00, 30.00, 32.00, 86.00, 93.00, 94.00, 95.00, 96.00, 97.00, 98.00, 99.00, 100.00, 101.00, 103.00, 163.00, 165.00, 167.00, 176.00, 177.00, 178.00, 179.00, 180.00, 181.00, 182.00, 183.00, 184.00, 185.00, 186.00, 187.00, 188.00, 189.00, 190.00, 191.00, 192.00, 193.00, 194.00, 195.00, 196.00, 197.00, 198.00, 199.00, 200.00, 201.00, 202.00, 203.00, 204.00, 205.00, 206.00, 207.00, 208.00, 209.00, 210.00, 211.00, 212.00, 213.00, 214.00, 215.00, 216.00, 217.00, 218.00, 219.00, 220.00, 221.00, 222.00, 224.00, 399.00, 328.00, 92.00, 366.00, 409.00, 410.00, 411.00, 412.00, 413.00, 414.00, 415.00, 416.00, 417.00, 418.00, 419.00, 420.00, 421.00, 422.00, 423.00, 424.00, 425.00, 426.00, 427.00, 428.00, 429.00, 430.00, 431.00, 432.00, 433.00, 434.00, 435.00, 436.00, 437.00, 438.00, 439.00, 440.00, 441.00, 442.00, 443.00, 444.00, 445.00, 446.00, 447.00, 448.00, 449.00, 450.00, 451.00, 452.00, 453.00, 454.00, 455.00, 456.00, 457.00, 458.00, 459.00, 460.00, 461.00, 462.00, 463.00, 464.00, 465.00, 466.00, 467.00, 468.00, 469.00, 470.00, 471.00, 472.00, 473.00, 474.00, 475.00, 476.00, 477.00, 478.00, 479.00, 480.00, 481.00, 482.00, 483.00, 484.00, 485.00, 486.00, 487.00, 488.00, 489.00, 490.00, 492.00, 639.00, 640.00, 641.00, 642.00, 644.00, 645.00, 647.00, 650.00, 651.00, 653.00, 331.00, 333.00, 511.00, 634.00, 374.00, 45.00, 47.00, 109.00, 579.00, 49.00, 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, 670.00, 672.00, 673.00, 674.00, 675.00, 676.00, 677.00, 678.00, 679.00, 680.00, 681.00, 682.00, 683.00, 684.00, 685.00, 686.00, 687.00, 688.00, 689.00, 690.00, 691.00, 692.00, 693.00, 694.00, 695.00, 696.00, 697.00, 698.00, 699.00, 700.00, 701.00, 702.00, 703.00, 704.00, 705.00, 706.00, 707.00, 708.00, 709.00, 710.00, 711.00, 712.00, 713.00, 714.00, 715.00, 716.00, 717.00, 718.00, 719.00, 720.00, 721.00, 722.00, 723.00, 724.00, 725.00, 726.00, 727.00, 728.00, 729.00, 730.00, 731.00, 732.00, 733.00, 734.00, 735.00, 736.00, 737.00, 738.00, 739.00, 740.00, 742.00, 906.00, 308.00, 55.00, 266.00, 268.00, 526.00, 528.00, 788.00, 789.00, 790.00, 791.00, 792.00, 908.00, 794.00, 905.00, 907.00	Superseded by Patch 921.01

Table 2–2: Summary of Base Operating System Patches

Patch IDs	Abstract
Patch 7.00 OSF510-037B	**Patch:** Threaded programs do not terminate **State:** Existing This patch fixes hangs in threaded programs with subprocesses created with nfork(NULL). Examining one of the hanging subprocesses shows that it has called fopen() and is waiting for the iobptr mutex in _findiop().
Patch 15.00 OSF510-009A	**Patch:** libst shared library fix **State:** Existing This patch fixes a problem with two routines in the libst library, st_obj_open() and st_obj_write(). The ability to change a file permission using these two libst routines is denied if a group has write permissions.
Patch 17.00 OSF510-009B	**Patch:** libst static library fix **State:** Existing This patch fixes a problem with two routines in the libst library, st_obj_open() and st_obj_write(). The ability to change a file permission using these two libst routines is denied if a group has write permissions.
Patch 19.00 OSF510-036	**Patch:** Fix for booting problem via network interface **State:** Existing This patch solves a problem which could prevent a V5.1 kernel from booting via a network interface. It corrects a timing issue which affects processors with speeds in excess of 700MHz.
Patch 38.00 OSF510-021	**Patch:** Fix for panic that occurs when kloadsrv is restarted **State:** Existing This patch fixes a system panic that may occur when /sbin/kloadsrv is restarted.
Patch 40.00 OSF510X11-001	**Patch:** Fix for lbxproxy utility **State:** Existing This patch fixes a problem where the X windows lbxproxy utility, which is used to make Low Bandwidth X (LBX) connections to an X server, did not accept local connections.
Patch 51.00 OSF510-043	**Patch:** Change to kloadsrv and hotswapd entries **State:** Existing This patch changes the kloadsrv and hotswapd entries in the /etc/inittab file. The change will prevent possible problems with dynamically loaded kernel modules when shutting down to single user mode.
Patch 74.00 OSF510-018	**Patch:** Fixes environmental warning in GS systems **State:** Existing This patch fixes a problem on the AlphaServer GS80, GS160, and GS320 platforms where the system will issue an environmental warning and shut itself down when it reaches a critical temperature, even though this temperature is safe for the power supply.
Patch 82.00 OSF510CDE-001	**Patch:** List of application groups is not re-created **State:** Existing This patch fixes a problem where the Common Desktop Environment (CDE) Application Manager did not re-create the list of application groups at login. After customizing the application groups, users would see the old groups instead of the new groups.

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

Patch 105.00 OSF510-017	**Patch:** Prevents not currently mounted warning messages **State:** New This patch prevents "not currently mounted" warning messages from being displayed for file systems the user did not request to umount.
Patch 107.00 OSF510X11-007	**Patch:** Fix for tcl **State:** New This patch fixes a problem in which tclhelp and any other tool using #!/usr/bin/wishx as the interpreter fails when additional versions of tcl are installed in /usr/local.
Patch 111.00 OSF510DX-003	**Patch:** Fix for smsd crash **State:** New This patch fixes intermittent crashes of the SysMan Station daemon (smsd) that are most likely to occur at system startup time, midnight, or during reconfiguration of system components. This crash would render a connected SysMan Station client unusable.
Patch 121.00 OSF510CDE-002	**Patch:** Fix for dtlogin **State:** New This patch fixes a problem where the Common Desktop Environment (CDE) login daemon, dtlogin, core dumps occasionally when servicing requests from XDMCP clients such as X terminals or PCs running X servers.
Patch 130.00 OSF510-067	**Patch:** Fix for lock hierarchy violation panic **State:** New. Supersedes patch OSF510-034 (80.00) This patch corrects the following: • Fixes a problem that can occur under certain circumstances with an IPv6 packet that contains a routing header. This could possibly crash a machine functioning as an IPv6 router. This was only reproduced with manually generated packets. • Under certain circumstances a Tru64 UNIX system configured with IPv6 can panic with a lock hierarchy violation. This panic can occur on any system running Tru64 UNIX with IPv6 enabled and configured.
Patch 150.00 OSF510-052	**Patch:** Fix for advscan **State:** New This patch fixes a problem where advscan -a -g does not display bootable partitions properly.
Patch 169.00 OSF510DX-009	**Patch:** Fix for bttape **State:** New bttape now uses the lock file /usr/run/bttape.pid for checking multiple instances. Also, the default addlist and fslist are created appropriately.
Patch 171.00 OSF510-098	**Patch:** Fix for voldctl stop command **State:** New This patch corrects the voldctl stop command behaviour for cluster support.
Patch 173.00 OSF510-057	**Patch:** fixso command causes segmentation fault **State:** New This patch fixes a problem with the /usr/ucb/fixso command that can cause a segmentation fault.

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

Patch 226.00 OSF510-107B	**Patch:** Fix for delayed AdvFS requests **State:** New This patch corrects some I/O rate fluctuations and thread unresponsiveness that had been seen when vm free pages dropped to a low level and used pages were being recycled.
Patch 232.00 OSF510-158B	**Patch:** Adds support for activating temporary data logging **State:** New This patch provides support for activating temporary data logging on a mount point.
Patch 245.00 OSF510-198	**Patch:** Install does not allow subset name with an underscore **State:** New. Supersedes patch OSF510-046 (9.00) This patch fixes the following problems: • Fixes a problem with the installation process rejecting a subset name with an underscore character on a V5.1 system. Specifically, when a user was trying to install the IBM MQSeries Documentation Base subset, MQS_HTML_PUBS. • Fixes a problem with the deletion process on a cluster system when a member node is running /usr/bin/csh. The process fails with a command not found error. • Fixes a problem with the deletion process not terminating when the C DELETE phase of the subset control program fails.
Patch 255.00 OSF510-090	**Patch:** tar -F ignores files named err **State:** New. Supersedes patch OSF510-164 (253.00) This patch corrects the following problems: • 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 • 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.
Patch 262.00 OSF510-074B	**Patch:** Fix for loader and ldd **State:** New This patch fixes the following problems: • Fixes a loader problem with rpaths on shared libraries, a loader problem when libraries loaded in -taso mode were loaded above the -taso address range, a problem detecting incorrectly specified _RLD_ARGS values, and a problem handling the RHF_BIND_NOW object file bit. • Fixes a problem with /usr/ucb/ldd. Previously the _RLD_ARGS environment variable was not recognized.

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

Patch 265.00 OSF510-205B	**Patch:** loader does not report error **State:** Supersedes patches OSF510-028 (78.00), OSF510-147 (263.00) This patch fixes the following problems: • Fixes a problem where applying spike to some binaries results in a 100% performance degradation. • Fixes a problem where spike may fail to delete the low instruction of a pair of related instructions, causing it to abort with a runtime error. • Fixes a problem that may cause the /usr/ucb/spike post-link optimization tool to crash. • Fixes a /sbin/loader problem that causes the ldr_inq_region() call to not report an error when an invalid region parameter is passed as a parameter to the call.
Patch 270.00 OSF510DX-017	**Patch:** Updates Netscape Communicator to Version 4.76 **State:** Supersedes patch OSF510DX-001 (44.00) This patch corrects the following problems: • Fixes a security vulnerability (called the Brown Orifice) in Netscape Communicator Version 4.72 by updating Netscape Communicator to Version 4.75. • Updates Netscape Communicator to Version 4.76 to fix missing default MIME types in Netscape Communicator 4.75.
Patch 283.00 OSF510-104	**Patch:** Security (SSRT0682U) **State:** New. Supersedes patch OSF510-096 (281.00) This patch corrects the following: • Fixes a problem in which rexecd fails to establish stderr. If the client rexec() function call specifies a secondary socket for stderr, connects to rexecd hang. • 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. Compaq has corrected this potential vulnerability.
Patch 285.00 OSF510-118	**Patch:** Fixes C++ runtime errors **State:** New This patch fixes C++ runtime errors.
Patch 302.00 OSF510DX-007	**Patch:** dop tool causes segmentation fault **State:** New This patch fixes a problem in which the dop tool would cause a segmentation fault when a non-root user entered the root password.
Patch 304.00 OSF510-137	**Patch:** Running cord on libraries causes infinite loop **State:** New This patch fixes an infinite loop that occurs when using cord on a library compiled with -g3. If the library has unused static routines that are optimized away, cord may go into an infinite loop.
Patch 306.00 OSF510-116	**Patch:** Fixes a C++ compiler error **State:** New This patch fixes a C++ compiler error.
Patch 312.00 OSF510-076	**Patch:** Fixes a problem of the ATM setup script failing **State:** New This patch fixes a problem of the ATM setup script failing when configuring an elan if the lane subsystem is not loaded.

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

Patch 318.00 OSF510-081	**Patch:** Fix for newgrp command **State:** New This patch corrects the problem where newgrp(1) fails if the file /etc/group contains multiple lines for one group.
Patch 320.00 OSF510DX-011	**Patch:** Fixes a problem in diskconfig **State:** New 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 324.00 OSF510-142B	**Patch:** Fix for libots3 **State:** New This patch fixes the following problem in the Compaq C compiler: • An optimizer problem that caused a failure in the llogin UNIX command. • An optimizer problem that caused incorrect run-time results for an OpenMP program. • A problem in the parallel processing support library that caused incorrect run-time results for an OpenMP program.
Patch 326.00 OSF510-082	**Patch:** Security (SSRT0672U) **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. Compaq has corrected this potential vulnerability.
Patch 335.00 OSF510-141	**Patch:** Fix for DVD file system problem **State:** New This patch addresses two issues with the DVD File system: • When directory entries are large enough to overflow a user's buffer and require multiple calls to complete, DVDFS fails because it does not properly calculate the continuation point for successive calls. • Logical block numbers are not properly calculated after the first directory data read.
Patch 337.00 OSF510DX-008	**Patch:** Fix for bindconfig **State:** New This patch fixes the problem of OutOfOrder hide stack trace, which occurs when an invalid domain name is entered during bindconfig.
Patch 339.00 OSF510CDE-005	**Patch:** Fix for dtpad **State:** New This patch fixes a problem where, if dtpad cannot allocate enough memory, it will exit and leave a zero-length file in place of the file being edited.
Patch 341.00 OSF510-197	**Patch:** Fix for ksh hang **State:** New This patch fixes a problem where the Korn shell (ksh) could hang if the user pastes a large number of commands to it when it is running in a terminal emulator window (such as an xterm).
Patch 343.00 OSF510-114	**Patch:** Fix for vi core dump **State:** New This patch fixes a problem in which the vi editor core dumps when it finds invalid syntax during a substitute operation.

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

Patch 345.00 OSF510-187	**Patch:** Cannot create builds with CAMDEBUG enabled **State:** New This patch fixes a problem of not being able to create builds with CAMDEBUG enabled.
Patch 349.00 OSF510-121	**Patch:** Corrects memory leak in XTI socket code **State:** New This patch corrects a memory leak in the XTI socket code.
Patch 351.00 OSF510-093	**Patch:** Fix for Turbolaser panic **State:** New This patch prevents a panic on TurboLaser systems with a DE600 in pci slot 0. Mis-identification of the DE600 in pci slot 0 causes data structure corruption. TurboLaser systems include the following: AlphaServer 8200 AlphaServer 8400 AlphaServer GS60 AlphaServer GS60E AlphaServer GS140 A DE600 is a single-port 10/100 Mbps Fast Ethernet NIC.
Patch 353.00 OSF510-190	**Patch:** Fix for fsx utility **State:** New This patch fixes a problem in which the fsx utility would not correctly handle the -s switch.
Patch 360.00 OSF510DX-012A	**Patch:** Nodes in cluster unable to set high temp threshold **State:** New This fix corrects a problem in which nodes in a cluster are unable to set their high temperature thresholds.
Patch 362.00 OSF510DX-012B	**Patch:** Cluster nodes unable to set high temp threshold **State:** New This fix corrects a problem in which nodes in a cluster are unable to set their high temperature thresholds.
Patch 364.00 OSF510-186	**Patch:** (SSRT1-15, SSRT0713U) **State:** 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. Compaq has corrected this potential vulnerability.
Patch 368.00 OSF510-086	**Patch:** rdist utility causes segmentation fault **State:** New This patch corrects a problem in the rdist utility which was causing segmentation faults on files with more than one link.
Patch 372.00 OSF510-127	**Patch:** Kernel memory fault occurs when using tablet **State:** New This patch fixes a kernel memory fault which occurs while using a tablet instead of a mouse.
Patch 385.00 OSF510-075	**Patch:** Fixes problem in exit status value of swapon utility **State:** Existing This patch fixes a bug in the exit status value of the swapon utility.
Patch 387.00 OSF510-143	**Patch:** CDFS media burned in 2001 shows the wrong dates **State:** New CDFS media burned in 2001 shows the wrong dates.

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

Patch 391.00 OSF510-128	**Patch:** System crash when accessing the FDI floppy **State:** New This patch corrects the following: • Compaq has determined in laboratory testing that there is a theoretical possibility that during read and write operations to the floppy disk on DS10, DS10L and ES40 AlphaServers and XP900 AlphaStations, a single byte of data may be inaccurately read or written without notice to the user or system. The potential for this anomaly exists only if floppy data read and write operations are attempted while there is extremely heavy traffic on these Alpha systems' internal input/output busses. Although Compaq has observed the anomaly only in laboratory tests designed to create atypical system stresses, including almost constant use of the floppy disk drive, we are supplying this patch to address this potential issue. • Corrects a potential system crash when accessing the FDI floppy.
Patch 397.00 OSF510-222	**Patch:** Fix for grep command **State:** Supersedes patch OSF510-031 (42.00) This patch fixes a problem with the grep command in which the options -p -v together do not produce any output.
Patch 496.00 OSF510-301	**Patch:** Fix for problems in Compaq C compiler **State:** New. Supersedes patches OSF510-016 (66.00), OSF510-142A (322.00) This patch fixes the following problems: • An optimizer problem that caused the wrong answer to be produced for a program involving tail recursion. • An optimizer problem that caused a runtime error when compiling gcc using -feedback. • An optimizer crash when compiling a program using -ieee and -tune ev6. • An optimizer problem that caused a failure in the llogin UNIX command. • An optimizer problem that caused incorrect run-time results for an OpenMP program. • A problem in the parallel processing support library that caused incorrect run-time results for an OpenMP program. • A compiler problem that caused a runtime 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 -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. • Fixes the following problem in the parallel processing support library (libots3): – A problem in the parallel processing support library that caused incorrect run-time results for an OpenMP program.

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

Patch 498.00 OSF510CDE-008A	**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. Compaq has corrected this potential vulnerability.
Patch 500.00 OSF510CDE-008B	**Patch:** Security (SSRT1-80U) **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. Compaq has corrected this potential vulnerability.
Patch 513.00 OSF510-317B	**Patch:** Security (SSRT1-40U, SSRT1-41U, SSRT1-42U, SSRT1-45U) **State:** New A potential security vulnerability has been discovere 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). Compaq has corrected this potential vulnerability.
Patch 515.00 OSF510-317C	**Patch:** Security (SSRT1-48U) **State:** New A potential security vulnerability has been discovere 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). Compaq has corrected this potential vulnerability.
Patch 517.00 OSF510-317D	**Patch:** Security (SSRT1-48U) **State:** New A potential security vulnerability has been discovere 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). Compaq has corrected this potential vulnerability.
Patch 519.00 OSF510-244	**Patch:** Installation process does not support alternate root **State:** New This patch fixes a problem with the installation process not supporting alternate root. When a subset is specified on the command line in non-interactive mode, its required subset is incorrectly referenced relative to the default root, not the alternate root.
Patch 523.00 OSF510-319	**Patch:** Fix for evmget command **State:** New 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 525.00 OSF510DX-020	**Patch:** Fix for dxarchiver core dump problem **State:** New This patch corrects a dxarchiver core dump problem. The core dump occurs when Clear button is clicked after archiving operation is complete.

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

Patch 530.00 OSF510-300A	**Patch:** Fixes POSIX message queue issues **State:** Supersedes patch OSF510-176A (401.00) This patch corrects the following: • Fixes POSIX message queue issues seen with mq_open() and other calls with messsaging. • Fixes a problem that mq_close of a message queue does not call the function p4_delete_entry to free up the resource. Thus, for a process that keeps using mq_open and mq_close, it will eventually run out of descriptors.
Patch 532.00 OSF510-300B	**Patch:** Fix for ERRNO EMFILE 24 error **State:** Supersedes patch OSF510-176C (405.00) This patch corrects the following: • Fixes POSIX message queue issues seen with mq_open() and other calls with messsaging. • Fixes a problem that mq_close of a message queue does not call the function p4_delete_entry to free up the resource. Thus, for a process that keeps using mq_open and mq_close, it will eventually run out of descriptors.
Patch 534.00 OSF510-300C	**Patch:** Fix for POSIX 4 message queue **State:** Supersedes patch OSF510-176B (403.00) This patch corrects the following: • Fixes POSIX message queue issues seen with mq_open() and other calls with messsaging. • Fixes a problem that mq_close of a message queue does not call the function p4_delete_entry to free up the resource. Thus, for a process that keeps using mq_open and mq_close, it will eventually run out of descriptors.
Patch 536.00 OSF510-318	**Patch:** Updates the emx driver to v2.02 **State:** Supersedes patch OSF510-166 (278.00) This patch corrects the following: • Fixes a problem where cascaded switches can hang the system at failover time. • Updates the emx driver to V2.02. – 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. – Can't grow list panic which can occur on large fabrics.
Patch 540.00 OSF510-250	**Patch:** OSF510-250 **State:** Supersedes patch OSF510-148 (292.00) 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 or privilege management. Compaq has corrected this potential vulnerability.

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

Patch 550.00 OSF510-359	**Patch:** JIB graphic card fix **State:** New 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 558.00 OSF510-256	**Patch:** Fixes a volrecover error **State:** New This patch fixes a volrecover error of "Cannot refetch volume" when volumes exist only in a non-rootdg diskgroup.
Patch 561.00 OSF510DX-018	**Patch:** Fixes a core dump problem in dxkerneltuner **State:** New. Supersedes patch OSF510DX-021 (559.00) This patch corrects the following: • Fixes a core dump problem in dxkerneltuner. The core dump occurs when you try to find an attribute (using Find Attributes option under the Options menu) that does not exist. • Fixes a core dump when the dxkerneltuner is used and the Select Subsystem button is pressed twice.
Patch 567.00 OSF510-260	**Patch:** Fixes a kernel memory fault in procfs.mod **State:** Supersedes patch OSF510-171 (316.00) This patch corrects the following: • Corrects a problem where attaching to a program with a debugger will cause periodic timers to be lost and will make the program hang. • Fixes a kernel memory fault in procfs.mod.
Patch 569.00 OSF510-344	**Patch:** Fixes a problem with the disklabel command **State:** New This patch fixes a problem with the disklabel command. Disklabel was displaying large unsigned values as negative numbers.
Patch 571.00 OSF510-219	**Patch:** Fix for mv command **State:** New This patch fixes a problem in which 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 said file and directory are on different file systems.
Patch 573.00 OSF510-336	**Patch:** Fixes a NetRAIN problem **State:** New This patch fixes a problem in NetRAIN. NetRAIN interface creation now fails if any of the requested standby interfaces do not exist.
Patch 575.00 OSF510DX-019	**Patch:** Allows the dxsetacl utility to delete access ACLs **State:** New This patch allows the dxsetacl utility to delete access ACLs.
Patch 577.00 OSF510-234	**Patch:** Fixes the consumption of excessive CPU cycles **State:** New This patch fixes the consumption of excessive CPU cycles caused by rshd when SIA is enabled.

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

Patch 583.00 OSF510-257	**Patch:** Kernel leaves cached pointers to ksm data structure **State:** Supersedes patch OSF510-020 (21.00) This patch corrects the following: • Corrects a stack overflow panic encountered during the startup of the system management deamon(smsd) on configurations with more than 255 devices. • Fixes a problem within the kernel that could leave cached pointers to a ksm data structure after the ksm instance was removed from the hierarchy. A kernel memory fault or data inconsistency could result.
Patch 585.00 OSF510-284	**Patch:** Fixes a problem with the keyboard driver **State:** New This patch fixes a problem where the keyboard driver takes too long probing for the keyboard when a keyboard is not connected.
Patch 587.00 OSF510-308	**Patch:** Security (SSRT0743U, 88914, SSRT0743U) **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. Compaq has corrected this potential vulnerability.
Patch 589.00 OSF510-270	**Patch:** Fix for sort command **State:** New This patch corrects the behavior of the sort(1) command which now checks for duplicates with -c, -u, and -k flags.
Patch 591.00 OSF510-215	**Patch:** Fixes several problems found in the KZPEA driver **State:** Supersedes patch OSF510-041 (57.00) This patch corrects the following: • Fixes a panic or a system hang which could occur on a DS20E with drives attached to the motherboard SCSI interface (Adaptec 7895 based) or on an Ultra3 KZPEA SCSI adapter. In addition to system hangs or panics on configurations using Memory Channel adapters some configurations have exhibited SCSI device problems. • Fixes several problems found in the KZPEA driver that could result in memory corruption, bus hangs, and system panics. This patch also includes binary error logging support in the driver.
Patch 593.00 OSF510-349	**Patch:** Fixes a potential race deadlock **State:** New This patch fixes a potential race deadlock between vclean/ufs_reclaim and quotaon/quotaoff, when quota is enabled.
Patch 595.00 OSF510-223	**Patch:** Fixes a bug that causes a panic due to software error **State:** New This patch fixes a bug that would cause a panic due to a software error that removed some functionality in system security.
Patch 598.00 OSF510X11-015A	**Patch:** Security (SSRT0638U) **State:** New. Supersedes patch OSF510X11-017A (596.00) This patch corrects the following: • 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. • Allows the dxsetacl utility to delete access ACLs.

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

Patch 600.00 OSF510X11-017B	**Patch:** Security (SSRT0638U) **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 root directory compromise via lpr using X11.
Patch 602.00 OSF510X11-015B	**Patch:** Fix for dxsetacl utility **State:** New This patch allows the dxsetacl utility to delete access ACLs.
Patch 607.00 OSF510-233	**Patch:** Fix for line printer problem **State:** New This patch fixes a loss of data with the parallel line printer driver. Without this patch data from a print job may get dropped if multiple jobs are sent to the line printers in rapid succession.
Patch 609.00 OSF510-238	**Patch:** Fix for cp command **State:** New This patch fixes a problem in which cp(1) and cat(1) produce different file sizes when reading from a tape device. The solution was to change the I/O buffer size of the cp command from 64K to 8K.
Patch 616.00 OSF510X11-016A	**Patch:** OSF510X11-016A **State:** Supersedes patches OSF510X11-004A (135.00), OSF510X11-008A (137.00) This patch corrects the following: • Fixes two memory leaks in the X Window System's X library (Xlib) that can occur when creating and destroying Motif List, Text, and TextField widgets. • Provides enhanced support for UTF-8 and UCS-4 locales. • Fixes a problem with libX11.a and libX11.so that might cause a core dump by failing to initialize some variables in some Xlib internal functions.
Patch 618.00 OSF510X11-016B	**Patch:** Fix for libX11.a and libX11.so core dump problem **State:** Supersedes patches OSF510X11-004B (138.00), OSF510X11-008B (140.00) This patch corrects the following: • Fixes two memory leaks in the X Window System's X library (Xlib) that can occur when creating and destroying Motif List, Text, and TextField widgets. • Provides enhanced support for UTF-8 and UCS-4 locales. • Fixes a problem with libX11.a and libX11.so that might cause a core dump by failing to initialize some variables in some Xlib internal functions.
Patch 620.00 OSF510-252	**Patch:** Security (SSRT1-85U) **State:** New A potential security vulnerability has been discovered where, under certain circumstances, system integrity may be compromised. xntpd contains a potential buffer overflow that may allow unauthorized access to bin privileges. Compaq has corrected this potential vulnerability.
Patch 622.00 OSF510-220	**Patch:** Fix for rerouting problem seen on a cluster **State:** New 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 626.00 OSF510-236	**Patch:** Fix for the i2c kernel module **State:** Supersedes patch OSF510-172 (376.00) This patch corrects the following: • Fixes DS10/DS20 performance problems introduced with the i2c driver by using thread blocking, rather than event_timeout() and DELAY(). • Fixes various inefficiencies in the i2c kernel module, and fixes a lock hierarchy violation that could be seen with the generic kernel attribute lockmode turned on.
Patch 628.00 OSF510-340	**Patch:** BPF default packet filter causes system panic **State:** New This patch corrects a problem which could result in a system panic on close() if the BPF default packet filter is in use.
Patch 630.00 OSF510-309	**Patch:** Fixes streams based drivers from failing **State:** Supersedes patch OSF510-077 (378.00) This patch corrects the following: • Fixes a problem in which the system may panic with the panic string "Unaligned kernel space access from kernel mode". • Fixes streams based drivers from failing DRV_GETHANDLE, due to non supported driver_handle.
Patch 632.00 OSF510X11-019	**Patch:** Fixes problems with X server X Image Extension **State:** New This patch fixes problems with the X server X Image Extension (XIE).
Patch 636.00 OSF510-273	**Patch:** Corrects problems with joind **State:** Supersedes patch OSF510-152 (389.00) This patch corrects the following: • Corrects a problem with joind which caused it to respond to certain client dhcp requests via the wrong port. • Fixes a problem where joind may fail to clean up its lock files in /var/join.
Patch 649.00 OSF510-261	**Patch:** Possible hang occurs with libaio and libaio_raw **State:** New This patch warns a user of a possible hang that can occur when a program is linked to both libaio and libaio_raw.

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

Patch 655.00
OSF510-465

Patch: Fix for vrestore problems

State: OSF510-013 (13.00), OSF510-161 (175.00), OSF510-285 (406.00), OSF510-242 (408.00)

This patch fixes the following problems:

- A previous patch caused incomplete restores.
- A warning message is displayed when the path for the first file in a group of hard links is created without using original protection codes and property lists.
- A warning message is displayed and vrestore aborts if it fails to malloc space for a property list.
- A message which had been inserted at the end of the message file had the wrong message category (this could cause messaging confusion).
- An uninitialized variable in the code that restores property lists could cause malloc failures, memory faults, an "error setting extended attributes" message, and infinite loops using the -l option.
- Corrupted property list information could cause an infinite loop.
- Fixes problems in the vdump command:
  - Failed to flag compressed extended attributes records that are split across a vdump BLOCK boundary.
  - Overrides the -D option when source path describes a root fileset (Note: If you want to back up quota files, you must not use the -D option.)
  - Corrects “Rewinding” message to avoid a segfault with Internationalized messages.
  - Prevents a core dump from vdump when your message length is greater than MAX_MSG_SIZE. This will be a very rare occurence.
  - Modifies vdump to forward space to next file only if a norewind tape was specified.
- Fixes problems in the vrestore command
  - 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).
  - The -Q option is added to vrestore to allow the user to request that quota files are ignored (thus avoiding the time it takes to process them).
- Prevents a core dump from vdump when your message length is greater than MAX_MSG_SIZE. This will be a very rare occurence.
- Modifies vdump to forward space to next file only if a norewind tape was specified.
- Fixes a problem in vrestore where, when restoring from a norewind tape, it incorrectly interprets a value and fails with an error message that looks similiar to the following:

  vrestore: unable to open save-set </dev/ntape/tape0c>; [0] Successful

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

Patch 671.00 OSF510-375	**Patch:** Hardware manager inaccurately reports the CPU speed **State:** Superseded patch 76.00 (OSF510-038) This patch corrects the following: • Fixes a problem where the hardware manager inaccurately reports the CPU speed. It reported a CPU speed that was one MHz less than the correct speed. • Fixes a potential security problem.
Patch 744.00 OSF510-512B	**Patch:** Installs DECthreads V3.18-138 **State:** Supersedes patches OSF510-039B (34.00), OSF510-212B (227.00), OSF510-206B (228.00), OSF510-109B (230.00) This patch fixes problems for threaded applications running on Tru64 UNIX V5.1: • Fixes a bug in the POSIX Threads Library for Tru64 UNIX V5.1 where a terminating thread did not properly clear an enabled floating point unit, causing invalid floating point state on the next thread that is run. • Fixes a bug in the POSIX Threads Library for Tru64 UNIX V5.1 that would result in a DECthreads error return of EINVAL from the pthread mutex API routines. This error would be seen only when the thread stack had been user defined/changed, specifically seen when using the user level context switching (ucontext) routines. • Fixes a bug in the POSIX Threads Library for Tru64 UNIX V5.1 that would result in a DECthreads Bugcheck and process termination. Threaded applications might encounter this problem when pthread_kill() is used on a thread that is marked as blocked in the kernel. • Installs DECthreads V3.18-138 which fixes problems that may affect threaded programs running on Tru64 UNIX V5.1.
Patch 746.00 OSF510-367A	**Patch:** Fix for problem in studio.h **State:** New This patch fixes a problem in <stdio.h> where the interface renaming conditionals for fgetpos() & 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 748.00 OSF510-367B	**Patch:** OSF510-367B **State:** New This patch fixes a problem in <stdio.h> where the interface renaming conditionals for fgetpos() & 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.

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

Patch 761.00 OSF510-400	**Patch:** Security (SSRT0689U, SSRT1-26, SSRT0788U, SSRT0781U) **State:** Supersedes patches OSF510-033 (1.00), OSF510-019 (2.00), OSF510-027 (3.00), OSF510-037A (5.00), OSF510-051 (87.00), OSF510-071 (88.00), OSF510-061 (90.00), OSF510-154 (233.00), OSF510-177 (234.00), OSF510-145 (235.00), OSF510-151 (236.00), OSF510-123 (237.00), OSF510-183 (238.00), OSF510-130 (239.00), OSF510-091 (240.00), OSF510-146 (241.00), OSF510-150 (243.00), OSF510-283 (501.00), OSF510-253 (502.00), OSF510-208 (504.00), OSF510-499 (749.00), OSF510-520 (750.00), OSF510-398 (751.00), OSF510-529 (752.00), OSF510-430 (753.00), OSF510-393 (754.00), OSF510-401 (755.00), OSF510-530 (756.00), OSF510-422 (757.00), OSF510-381 (758.00), OSF510-467 (759.00) This patch corrects the following problems: • Fixes a problem of the getaddrinfo() library call returning a failing status. • Increases the number of places of precision for formatted printing of long doubles. • Fixes the problem that, on rare occasions, the C runtime library atof() and strtod() functions (and other functions that may use them) may produce an incorrect result. The error would only be in the least significant digit of the mantissa (a rounding error). • Fixes hangs in threaded programs with subprocesses created with nfork(NULL). Examining one of the hanging subprocesses shows that it has called fopen() and is waiting for the iobptr mutex in _findiop(). • Fixes the printing of 0.0 when precision is specified for a %g type conversion. • Fixes a problem where a TZ environment variable setting of ":" yields incorrect (or missing) time zone information after calling tzset() and incorrect error reporting from mktime(). • 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. Compaq has corrected this potential vulnerability. • Fixes a performance problem with freeing memory in threaded applications, when many allocations of the same size have been made. It also fixes a problem when the __sbrk_override malloc tuning variable is set which caused malloc to try to allocate too much memory. • Fixes a problem with the mallinfo() call which can cause an application to fail if run on a RAD other than 0. • Fixes a problem with the mallinfo() call which can cause an application to fail if run on a RAD other than '0'. • Restores correct behavior that existed on pre-V5.0 releases for ecvt() and fcvt(). Floating point exceptions and core dumps no longer occur when denormalized values are passed to ecvt() and fcvt(). • Resolves issues with customer applications that experienced floating point exceptions and core dumps when passing denormalized values to ecvt() and fcvt() that subsequently caused INFORMIX databases to crash. • Fixes the return values for vwprintf() functionality when used with wide characters. • Increases the input buffer size limits for the scanf family of functions to the MAXINT input buffer size.

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

Patch 761.00 continued	• Fixes the problem of optimized programs printing incorrect values for long doubles. • Adds logic that implements maximum size checks for input width descriptors on numeric scanf() format elements. • Corrects a regular expression performance problem in libc. • Fixes a potential online help build problem when dthelptag is used to compile online help files in a multibyte locale. • Fixes regular expression handling with non-default locale settings. • Fixes a regular expression matching problem in multibyte locales. • Corrects a problem in which the rsh [host] -l [user] [command] command returns "permission denied". • 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. Compaq has corrected this potential vulnerability. • Fixes a problem with fopen. fopen was returning "file not found" when there was insufficient memory available to allocate the FILE structure. fopen now returns "not enough space" for this case. • Fixes a problem with strerror where buffers could not be allocated. • Fixes a problem in which the RPC TCP server incorrectly tries to write to a socket that has already been closed by a client. • 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. Compaq has corrected this potential vulnerability. • Fixes a regular expression problem with the REG_NEWLINE flag of the regexec() routine. • Fixes confusing prompts for synchronized password updates when only one named mechanism is listed in the passwd command dialog box. • Fixes a problem in fread() where excessive I/O was taking place for large amounts of data, causing performance problems. It also addresses a failure in fread() to properly handle data sizes that have representations greater than 32 bits (2^32 of data). • Fixes a segmentation fault problem with long LOCPATH and LANG values. • Fixes a problem in mktime() when adjusting for a tm struct containing an invalid tm_isdst (Daylight Savings Time) setting. • Fixes a regular expression performance problem as well as two bugs that posed potential regular expression problems for multibyte locales • Fixes an application core dump problem when the LANG environment variable is too long. • Fixes a problem in fwrite() where it was failing when the total number of bytes to be written is larger than 2 GB.

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

Patch 766.00 OSF510CDE-010A	**Patch:** Security (SSRT0788U, SSRT0753U, SSRT0752U, SSRT0757U) **State:** Supersedes patches OSF510CDE-006 (310.00), OSF510CDE-009 (565.00), OSF510CDE-003 (158.00), OSF510CDE-013A (762.00), OSF510CDE-015A (763.00), OSF510CDE-014 (764.00) This patch corrects the following: • Fixes a problem on multi-head systems in which the unlock display only works if the default display is screen 0. • Fixes the problem of palette files not being read from /etc/dt/palettes. • Fixes a dtmail problem in which a From line with quotes in it incorrectly finds the date of the mail message. This error is displayed on the main screen under the header Date and Time and shows up as Dec. 31 or as a blank field. • 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. Compaq has corrected this potential vulnerability. • Fixes the dtprintinfo memory fault problem with a long LANG value. • Fixes a potential security vulnerability in CDE Subprocess Control Service(dtspcd). • A potential security vulnerability has been discovered where, under certain circumstances, system integrity may be compromised.This maybe in the form of large values of ENVIRONMENT variables and command line arguments. Compaq has corrected this potential vulnerability.
Patch 770.00 OSF510CDE-010B	**Patch:** Security (SSRT0788U, SSRT0757U, SSRT0753U, SSRT0752U) **State:** New. Supersedes patches OSF510CDE-013B (767.00), OSF510CDE-015B (768.00) This patch corrects the following: • 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. Compaq has corrected this potential vulnerability. • Fixes the dtprintinfo memory fault problem with a long LANG value. • A potential security vulnerability has been discovered where, under certain circumstances, system integrity may be compromised.This maybe in the form of large values of ENVIRONMENT variables and command line arguments. Compaq has corrected this potential vulnerability.
Patch 772.00 OSF510DX-023	**Patch:** Fix for dxproctuner utility **State:** New 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.
Patch 774.00 OSF510-429	**Patch:** Fixes a problem with siacfg **State:** New This patch fixes a problem with siacfg. The siacfg -A option was not working. Specifically, the BSD mechanism is still the last mechanism listed. The restriction has been removed.

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

Patch 777.00 OSF510X11-024	**Patch:** Fixes window corruption on Oxygen VX1 graphics card **State:** Supersedes patches OSF510X11-010 (246.00), OSF510X11-013 (247.00), OSF510X11-014 (249.00), OSF510X11-003 (119.00), OSF510X11-012 (287.00), OSF510X11-018 (505.00), OSF510X11-020 (506.00), OSF510X11-023 (507.00), OSF510X11-022 (509.00), OSF510X11-029 (775.00) This patch fixes the following problems: • Fixes a memory leak in the X server that could occur when a client repeatedly created and destroyed buffers for the X Window System Multibuffering Extension (XmbufCreateBuffers/XmbufDestroyBuffers). • Fixes a problem where the X server does not display windows properly for the 128th and subsequent clients. • Changes the X server to dynamically retrieve its vendor string information when running on COSIX64. • Provides the Xserver library for a new graphics card. • Corrects blocks of erroneous pixels left behind when dragging CDE "pplication manager icons on the desktop. • 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 an Xserver crash when using the GTK on systems using the Oxygen VX1 graphics card. • Fixes the Xserver problem where, when PanoramiX is enabled and using CDE, icons from dtfile can not be seen on other than the left screen while being moved. • Fixes a problem that can cause CDE pop-up menus to appear on the wrong screen when you are running a multi-head system with the PanoramiX extension enabled. • 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 window corruption on Oxygen VX1 graphics card if backing store/save unders are enabled.
Patch 782.00 OSF510-370	**Patch:** Fixes a problem with hwautoconfig **State:** New This patch fixes a problem with hwautoconfig, which was causing the I20 management driver to crash the system with a Kernel Memory Fault when it was loaded.

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

Patch 785.00 OSF510-371	**Patch:** Fix for V5.1 dynamic loader **State:** Supersedes patches OSF510-005 (64.00), OSF510-113 (256.00), OSF510-105 (257.00), OSF510-205A (258.00), OSF510-074A (260.00), OSF510-221 (521.00), OSF510-495 (783.00) This patch fixes the following problems with the V5.1 dynamic loader: • Allows the loader to properly ignore unreferenced symbols when loading a shared library with a dlopen call. • Allows the loader to properly ignore loading a library with the correct library name but an incorrect library version. • Fixes an /sbin/loader problem dealing with absolute value symbols when their value was -1. • Fixes a problem in the /sbin/loader dynamic loader that can cause a crash. It also fixes a problem with the output for the ldd command, where the output was always going to stderr rather than stdout. • Fixes a problem that may cause the /usr/ucb/spike post-link optimization tool to crash. • A /sbin/loader problem that causes the ldr_inq_region() call to not report an error when an invalid region parameter is passed as a parameter to the call. • Fixes a loader problem with rpaths on shared libraries, a loader problem when libraries loaded in -taso mode were loaded above the -taso address range, a problem detecting incorrectly specified _RLD_ARGS values, and a problem handling the RHF_BIND_NOW object file bit. • Fixes a problem with /usr/ucb/ldd. Previously the _RLD_ARGS environment variable was not recognized. • Fixes a problem in /sbin/loader. It corrects certain loader failures reported for mismatched shared library versions. • Fixes a loader core dump that occurs when invoking certain call_shared executables that have been processed by post-link instrumentation tools. • Fixes the -ignore_all_versions and -ignore_version flags for the run-time loader (/sbin/loader).
Patch 787.00 OSF510-531	**Patch:** Fix for lsmsa product **State:** New This patch addresses a problem in the display of disk controller disk hierarchy by the lsmsa product.

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

Patch 796.00 OSF510X11-026	**Patch:** Fixes a memory leak in the X server **State:** Supersedes patches OSF510X11-002 (70.00), OSF510X11-011 (280.00) This patch fixes the following problems: • Fixes a problem on systems with a PowerStorm 4D10T (ELSA Gloria Synergy, SN-PBXGK-BB) graphics card or a PCI To Ethernet/Graphics Combo Adapter (3X-DEPVD-AA). Sometimes lines and images are not drawn correctly in scrolled windows. • Fixes synchronization and drawing problems in the X server for the PowerStorm 4D10T (ELSA Gloria Synergy, SN-PBXGK-BB) graphics card. • Fixes a memory leak in the X server on systems with a PowerStorm 4D10T (ELSA Gloria Synergy, SN-PBXGK-BB) graphics card that could occur when a client repeatedly created and destroyed buffers for the X Window System Multibuffering Extension (XmbufCreateBuffers/XmbufDestroyBuffers). • The Elsa Gloria Comet card does not correctly draw nested shaded boxes or anything similar.
Patch 798.00 OSF510-432	**Patch:** Fixes a problem in uucp **State:** New 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 800.00 OSF510-412	**Patch:** Fixes problems in the Tru64 UNIX Assembler **State:** Supersedes patches OSF510-132 (288.00), OSF510-103 (290.00), OSF510-274 (538.00) This patch corrects the following: • When assembling a .s file containing a data declaration directive (such as .byte) that specifies a list of values greater than 74, a fatal "yacc stack overflow" condition is raised. • A main procedure's prologue description will overwrite that of an alternate entry point when they both share the same address and they both specify their own .prologue directive. • A .s file that contains .align directives in its text section that is assembled at an optimization level greater than O0 may produce a series of zeros in its text section which, if executed, would cause the program to halt. • The -arch and -tune command line switches were essentially being ignored. • Code generated by the assembler for emulated ldb/ldbu/ldw/ldwu instructions produces incorrect results leading to a linker optimization that produces an invalid executable. • Code generated by the assembler for emulated ldb/ldbu/ldw/ldwu instructions produces incorrect results leading to a linker optimization that produces an invalid executable. • Code generated for loads with offsets larger than 32K is incorrect. • Incorrect addresses are generated when symbolic arithmetic is used, and when the address in question extends beyond the intitial 64K boundary of a section. • A prodecure with no instructions causes the assembler to segfault. • A prodecure with no instructions causes line number generation to segfault. • Data declared using the .gprel32 directive was not being longword aligned. • The relocation count for a program that contains a section that has in excess of 65535 reloctions will be incorrect, resulting in a bad link and an invalid executable. • An entry (PDSC_FLAGS_BASE_REG_IS_FP) was not being set correctly in a short-form stack-frame RPD when a .frame directive specified register 15. • When two entry points to a procedure (main or alternate) share the same address, the assembler generates four nop profiling instruction sequences for each one when the -pg switch is specified. This causes post-link tools, such as spike, problems. • When a main and an alternate entry point share both an address and a prologue, the assembler associates the prologue with the alternate entry and not the main, resulting in the assembler not generating an RPD because it does not see a prologue for the main entry. • The assembler miscalculates the number of relocations present in the .text section if a jmp/jsr instruction was specified without a symbol as an operand. This can result in a linker error.

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

Patch 800.00 continued

- Resolves four incompatibilities between the new (as of V5.1) and old assemblers that are needed to support a future port of gcc to Tru64 UNIX. These changes are included in this version (3.04.33) of the assembler:
  - The assembler has never generated a section header for zero-sized sections, or a symbol table entry for a label symbol that is associated with such a section. This essentially correct behavior represents an incompatibility with the old assembler and has been changed with this patch.
  - The assembler was not including symbols for numeric constant label symbols in the symbol table. It is now.
  - The assembler can produce incorrect scoping for local symbols, resulting in incorrect association of symbols with their containing procedures.
  - The assembler's association of label symbols to their files of origin was incorrect in certain circumstances:

```
        .file   1 "file1.cxx"
gcc2_compiled.:
__gnu_compiled_cplusplus:
        .file   2 "file2.h"
        .file   3 "file3.h"
        .text
```

    In this example, label1 and label2 are mistakenly associated with file3.h due to the assembler's practice of establishing file context based on the instruction with which a given label was associated, which in this case was the first intruction in the .text section. File context is positional and in this case both labels should be associated with file1.
- Shipped as Version 3.04.34 of the Tru64 UNIX Assembler, this patch resolves three assembler problems:
  - This unusual case takes the following combination of factors:
    - ☐ A label defined at the head of the .rdata section.
    - ☐) Multiple file references (use of the .file and .loc directives) such that the initial .rdata label and the entry label of the last procedure in the .text section are associated with different files.

    The resulting symbol table scoping information is invalid, and causes om to seg fault. Note that this symbol table, although incorrect, does not stop the object file from linking and executing properly if -om is removed from the mix.
  - The assembler improperly reorders an instruction which restores the stack pointer when assembling with optimization active. The scheduler has specific logic which prevents an addq or an lda instruction which restores sp (and is followed by a ret instruction) from being moved. The problem occurs in a case where a bis instruction is used in order to restore sp. This instruction is being reordered by the assembler and it must not be.
  - The assembler is not adding a terminating NULL to the string specified as the argument for a .ident directive when the string is written to the object file. This causes the what command to produce incorrect return values.

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

Patch 803.00 OSF510DX-025	**Patch:** Fix for problem with dxaccounts **State:** OSF510DX-002 (36.00), OSF510DX-005 (122.00), OSF510DX-006 (124.00), OSF510DX-034 (801.00) This patch fixes the following dxaccounts problems: • A system running ASU experiences a dxaccounts crash problem when a user is deleted from PC User view. • The dxaccounts dialog messages are incorrectly displayed when a user is added with no password entry. • The dxaccounts utitlity is unable to create a new user from the PC Users view on a system with ASU installed. • The following problems can occur with the dxaccounts application on ASU system: – The dxaccounts utility crashes when the root icon is double clicked. – The full name of a new PC account is not mapped to a UNIX user. – Erasing a PC account's fields does not work; the values erased remain. – The default values of Home Directory, Login Script, and User Profile Path for a PC user are invalid. • Changing root's login/uid is enabled via cli/dxaccounts utilities. • Incorrect results of usermod -G. • The -x account_inactive | account_expiration options do not set the attributes. • Fixes a problem where the new home directory for a new user ID is created with the date and time stamp of the /usr/skel directory. • Fixes message fragments to make them I18N compatible. • 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. Compaq has corrected this potential vulnerability. • A core dump occurs when /etc/shells 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.
Patch 805.00 OSF510-460	**Patch:** Fix for kernel panic caused by ACL problem **State:** New This patch fixes the following problems: • 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 807.00 OSF510DX-035	**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. Compaq has corrected this potential vulnerability.

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

Patch 810.00 OSF510-510A	**Patch:** Fix for EVM core dump problem **State:** Supersedes patches OSF510-160A (294.00), OSF510-456 (808.00) This patch corrects the following: • Fixes a problem with the EVM daemon, evmd, where it will crash if /etc/rc.config contains a blank line. • Fixes a problem in which EVM subscribing clients, including /usr/sbin/evmlogger, will unexpectedly drop the connection to the EVM daemon. • Resolves an issue which can cause an Event Manager (EVM) client or the EVM daemon to core dump under rare circumstances.
Patch 812.00 OSF510-510B	**Patch:** EVM daemon may core dump **State:** Supersedes patch OSF510-160B (296.00) This patch corrects the following: • Fixes a problem with the EVM daemon, evmd, where it will crash if /etc/rc.config contains a blank line. • Resolves an issue which can cause an Event Manager (EVM) client or the EVM daemon to core dump under rare circumstances.
Patch 816.00 OSF510-515	**Patch:** Fixes a kernel memory fault panic **State:** Supersedes patches OSF510-004 (128.00), OSF510-210 (298.00), OSF510-217 (541.00), OSF510-286 (543.00), OSF510-417 (813.00), OSF510-516 (814.00) This patch corrects the following problems: • Fixes a cross RAD I/O hang problem with the ITPSA controller. • Fixes a problem that can cause a simple lock timeout or a kernel memory fault on EV6 systems using the ITPSA driver. • Fixes panics associated when multiple KZPCA-AA and/or KZPCM-AA host bus adapters are in the system. The expected panic string is "sc ws remove: SZ_IN_USE NOT set". • Fixes kernel memory faults, and/or I/O hangs with systems that have KZPCM-AA and/or KZPCA-AA. These errors can occur during large data transfers. • Fixes a panic in the ITPSA driver. It is seen when an abort to the SCSI rewind command is issued to a TLZ10 tape device. • Fixes a panic caused by SCSI bus resets with KZPCA HBAs. • Fixes a kernel memory fault panic during boot process while probing SCSI bus. • Fixes a kernel memory fault panic after an "ITPSA: itpsa_action - error converting path ID to ITPSA softc structure" message.
Patch 818.00 OSF510DX-033	**Patch:** Security (SSRT1-40U, SSRT1-41U, SSRT1-42U, SSRT1-45U) **State:** New. Supersedes patches OSF510DX-029A (859.00), OSF510DX-038 (778.00), OSF510DX-036 (780.00) 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. Compaq has corrected this potential vulnerability.

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

Patch 820.00 OSF510-450	**Patch:** Fixes a problem with the JJ printcap parameter **State:** Supersedes patches OSF510-080 (300.00), OSF510-245 (548.00) This patch corrects the following: • Introduces the JJ /etc/printcap parameter, which allows the user to choose either one TCP/IP connection for all jobs in the print queue (JJ=1), or a TCP/IP connection for each job in the print queue ( ). It also closes a timing hole that existed when lpd was shutting down. • Fixes a problem in which lpd hangs when printing to advanced server queues (using /dev/null). • Fixes a problem with the JJ /etc/printcap parameter.
Patch 823.00 OSF510-416	**Patch:** Fix for collect command **State:** Supersedes patches OSF510-026 (134.00), OSF510-247 (552.00), OSF510-433 (821.00) This patch corrects the following: • Fixes several problems with the collect command and it adds sysloging when collect suspends, resumes, or receives a signal. • Fixes collect's collector (/usr/sbin/collect) to correctly report the network interface load percentage. • Allows collect to filter out unnecessary file systems. • Adds support for the Mylex RAID controller as well as fixes several problems with the collect utility.
Patch 825.00 OSF510-458A	**Patch:** Fix for class scheduler **State:** Supersedes patch OSF510-280A (554.00) This patch corrects the following: • 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 827.00 OSF510-458B	**Patch:** Fix for class scheduler problem **State:** Supersedes patch OSF510-280B (556.00) This patch corrects the following: • 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 829.00 OSF510-486	**Patch:** Fix for verify command **State:** New This patch avoids core dumps in the verify command.

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

Patch 833.00 OSF510-543	**Patch:** Fix for showfdmn and rmvol programs **State:** Supersedes patches OSF510-230B (494.00), OSF510-485 (830.00), OSF510-455 (831.00) This patch corrects the following: • This AdvFS correction makes the balance and rmvol programs more interruptible by supplying a new option (-i). It also avoids wasting extent map entries and avoids a kmf in overlay_xtnt_map. • Balance was terminating before balancing the whole domain when the domain was very large (>4GB). • Fixes a potential problem with vdf and showfdmn, where they could incorrect display the message “showfdmn: No such file or directory”. • 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
Patch 835.00 OSF510X11-025A	**Patch:** Security (SSRT0753U) **State:** Supersedes patch OSF510X11-006A (142.00) This patch corrects the following: • Fixes various memory leaks in the Motif library (libXm) that could occur when creating and destroying Motif List, Text, and TextField widgets. • 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. Compaq has corrected this potential vulnerability.
Patch 837.00 OSF510X11-025B	**Patch:** Security (SSRT0753U) **State:** Supersedes patch OSF510X11-006B (144.00) This patch corrects the following: • Fixes various memory leaks in the Motif library (libXm) that could occur when creating and destroying Motif List, Text, and TextField widgets. • 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. Compaq has corrected this potential vulnerability.
Patch 840.00 OSF510DX-024	**Patch:** Security (SSRT1-40U, SSRT1-41U, SSRT1-42U, SSRT1-45U) **State:** New. Supersedes patch OSF510DX-026 (838.00) This patch corrects the following: • 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. Compaq has corrected this potential vulnerability. • Implements message catalog cleanup.

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

Patch 842.00 OSF510-534	**Patch:** Fix for rdump utility **State:** Supersedes patches OSF510-003 (68.00), OSF510-289 (563.00) This patch corrects the following: • Fixes a problem where a user could not dump to a regular file. • Fixes a core dump caused by using the rdump utility to back up data. • Now rdump command dumps data properly onto remote tape devices without receiving the signal SIGSEGV and dumping core.
Patch 844.00 OSF510-501	**Patch:** Security (SSRT0708U) **State:** Supersedes patch OSF510-138 (314.00) This patch fixes the following /usr/sbin/inetd problems: • A potential security vulnerability has been discovered, where, under certain circumstances, system integrity may be compromised. This may be in the form of inetd child process core dumping or failing to service incoming connection requests. Compaq has corrected this potential vulnerability. • inetd can terminate without notice and without a core file. • The disable keyword is being ignored when used in the /etc/inetd.conf.local configuration file. • The -h option does not restart any inetd children to handle requests because the parent still thinks one is running. • Allows the socket listen backlog in inetd(8) to be settable by command line option using the -l switch.
Patch 846.00 OSF510DX-039	**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. Compaq has corrected this potential vulnerability.
Patch 848.00 OSF510-407	**Patch:** SNMP returns wrong OID **State:** New This patch corrects the problem where SNMP getnext returns back the value of the wrong OID on queries in the FDDI MIB of os_mibs.
Patch 850.00 OSF510-427	**Patch:** Corrects a problem with a mirrored LSM volume **State:** New 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.
Patch 852.00 OSF510DX-037	**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. Compaq has corrected this potential vulnerability.
Patch 855.00 OSF510-394	**Patch:** KMF caused by malformed IPv6-in-IPv4 packets **State:** New. Supersedes patch OSF510-475 (853.00) This patch fixes a kernel memory fault caused by malformed IPv4-in-IPv4 packets.

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

Patch 857.00 OSF510DX-030	**Patch:** Security (SSRT1-40U, SSRT1-41U, SSRT1-42U, SSRT1-45U **State:** Supersedes patch OSF510DX-022 (581.00) This patch corrects the following: • 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. Compaq has 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 improper file access. Compaq has corrected this potential vulnerability.
Patch 861.00 OSF510DX-029B	**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. Compaq has corrected this potential vulnerability.
Patch 863.00 OSF510DX-040	**Patch:** Security (SSRT1-40U, SSRT1-41U, SSRT1-42U, SSRT1-45U) **State:** Supersedes patch OSF510DX-004 (162.00) This patch corrects the following: • Fixes a problem that was causing diskconfig to issue the error message "can't read "tminor": no such variable" upon startup. • 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. Compaq has corrected this potential vulnerability.
Patch 866.00 OSF510CDE-017	**Patch:** Security (SSRT1-40U, SSRT1-41U, SSRT1-42U, SSRT1-45U) **State:** New. Supersedes patch OSF510CDE-012 (864.00) This patch corrects the following: • 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. Compaq has 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 improper file access. Compaq has corrected this potential vulnerability.
Patch 868.00 OSF510DX-041	**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. Compaq has corrected this potential vulnerability.
Patch 870.00 OSF510CDE-011	**Patch:** Security (SSRT0767U) **State:** New A potential security vulnerability has been discovered where, under certain circumstances, system integrity may be compromised. The ttdbserverd contains a potential buffer overflow that may allow unauthorized access. Compaq has corrected this potential vulnerability.

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

Patch 872.00 OSF510X11-028A	**Patch:** Fix for X Toolkit library **State:** Supersedes patches OSF510X11-005A (112.00), OSF510X11-009A (114.00) This patch corrects the following: • Fixes a memory leak in the X Window System's X Toolkit library (Xt) that could occur when creating and destroying Motif List, Text, and TextField widgets. • Fixes the problem that XmStringGetLtoR() fails in dxhanziim when it runs in a C/en_US.ISO8859-1 locale. • Fixes a problem in the X Toolkit library (Xt) which could cause the TeMIP Iconic_map Presentation Module application (mcc_iconic_map) to crash.
Patch 874.00 OSF510X11-02O8B	**Patch:** Fixes for crash in X Toolkit library **State:** Supersedes patches OSF510X11-005B (115.00), OSF510X11-009B (117.00) This patch corrects the following: • Fixes a memory leak in the X Window System's X Toolkit library (Xt) that could occur when creating and destroying Motif List, Text, and TextField widgets. • Fixes the problem that XmStringGetLtoR() fails in dxhanziim when it runs in a C/en_US.ISO8859-1 locale. • Fixes a problem in the X Toolkit library (Xt) which could cause the TeMIP Iconic_map Presentation Module application (mcc_iconic_map) to crash.
Patch 876.00 OSF510-464	**Patch:** Fix for Event Manager channel monitoring function **State:** New This patch fixes a problem in which the Event Manager's channel monitoring function is temporarily disabled if the evmreload command is run.
Patch 878.00 OSF510-426	**Patch:** Fixes an ATM signalling problem **State:** Supersedes patch OSF510-079 (347.00) This patch corrects the following: • fFixes a problem of ATM signalling going into the "connection released" after a system reboot. • Fixes an ATM signaling problem.
Patch 880.00 OSF510-477	**Patch:** Fix for RIS/DMS serving in a TruCluster **State:** New This patch corrects the following: • A panic caused by a known problem, using a cluster as a RIS server. • A fix to RIS/DMS serving in a TruCluster.

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

Patch 882.00 OSF510-399	**Patch:** Fix for ld -f command **State:** Supersedes patches OSF510-022 (72.00), OSF510-153 (354.00), OSF510-108 (355.00), OSF510-120 (356.00), OSF510-102 (358.00), OSF510-258 (603.00), OSF510-240 (605.00) This patch corrects the following: • Fixes a spike problem. The problem results in an assertion and core dump when trying to spike a kernel. This patch is only needed if the post-link tool spike will be used on the Tru64 UNIX kernel. • Fixes a problem where the linker defined symbol _fpdata would end up being undefined if it was referenced by a program but not used by the linker. • Fixes link errors encountered when linking with -A. • Fixes two problems in the linker where it would erroneously report "multiply defined symbol" errors or "unresolved symbol" errors. – Modifies the linker's symbol resolution to enable it to recognize when a reference to a symbol defined in a shared library is replaced by a symbol defined in an object file or archive. – Modifies the linker to cause it to rescan shared libraries before reporting unresolved symbols. • Fixes two errors that occur when using the -f switch with the linker (ld): – Using the -f switch produces link errors. – Any unsupported switch beginning with -f gets interpreted to mean -f. • Fixes a potential optimization problem with the linker (/bin/ld). • Fixes a linker failure that can occur when linking a -non_shared executable with libexc.a. • The linker (/bin/ld) may corrupt the shared object registry file when -update_registry is specified with concurrent links.
Patch 884.00 OSF510-411	**Patch:** Addition of a tunable to the dli subsystem **State:** New This patch adds a tunable to the dli subsystem (mopsys_id) that provides the ability to disable the MOPSYS ID messages. These messages are sent every 10 minutes to inform bridges and routers of the system. They are on by default.
Patch 886.00 OSF510-395	**Patch:** Fix for kernel memory fault when using ATM **State:** Supersedes patches OSF510-056 (126.00), OSF510-338 (610.00), OSF510-312 (612.00) This patch corrects the following: • When running ATM Lan Emulation, using more than 4 ATM NetRAIN interfaces can result in recursive calls causing a “kernel stack not valid” halt. • Corrects a problem which could result in ATM/lane connection requests being dropped. • Fixes a kernel memory fault when using ATM.

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

Patch 888.00 OSF510-493	**Patch:** Fixes fixfdmn utility problem **State:** Supersedes patch OSF510-265 (614.00) This patch provides support for the /sbin/advfs/fixfdmn utility. The /sbin/advfs/fixfdmn utility is a tool that is used to check and repair corrupted AdvFS domains. Fixfdmn 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.
Patch 890.00 OSF510-489	**Patch:** Security (SSRT0664U) **State:** Supersedes patches OSF510-100 (379.00), OSF510-131 (381.00) This patch corrects the following: • 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. Compaq has corrected this potential vulnerability. • Corrects ftp daemon failure when using globbing string of several asterisks. Also contains additional corrections for the help command and character drop with the put command.
Patch 892.00 OSF510-408	**Patch:** Log in requests hang when enhanced security enabled **State:** Supersedes patches OSF510-054 (145.00), OSF510-055 (146.00), OSF510-072 (148.00), OSF510-170 (370.00), OSF510-254 (624.00) This patch corrects the following: • Corrects a problem in an Enhanced Security configuration where, at login time, if it is determined an account's password has expired, the "Old password:" prompt did not appear. Rather, the user is immediately prompted for their new password options and is allowed to change to a new password. This patch also allows a user logged into a system configured as a NIS client with Enhanced Security installed to change their password. • Fixes a problem in an Enhanced Security configuration. This patch restores the capability of being able to su to a user as root without being prompted or having to know the users password. • Fixes a problem for Enhanced Security configurations where the Maximum Login Interval (u_max_login_intvl) field was being ignored for account templates. • Fixes problems with the prpasswdd daemon hanging when there are numerous background processes simultaneously attempting to authenticate users to the system in an Enhanced Security environment. • Fixes a problem in which login requests can hang when enhanced security is enabled. • Fixes a problem where logins appear to be hung on standalone systems with Enhanced Security enabled.
Patch 896.00 OSF510DX-027	**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. Compaq has corrected this potential vulnerability.

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

Patch 898.00 OSF510-452	**Patch:** Fix for zdump utility **State:** New This patch corrects the following: • Fixes a problem in the zdump utility when time zone file names are specified as arguments without leading colons (:)s. • Fixes a regression in the -v output to display the current time.
Patch 900.00 OSF510-523	**Patch:** Fixes LSM configuration daemon (vold) **State:** New This patch is required to have CLSM clone rejection code to work properly.
Patch 902.00 OSF510CDE-016	**Patch:** Fix for window manager memory leak **State:** Supersedes patches OSF510CDE-004 (160.00), OSF510CDE-007 (393.00) This patch corrects the following: • Fixes a problem where the Window Manager (dtwm) intermittently hangs on a system which uses multiple displays. • Fixes a problem where the Common Desktop Environment (CDE) window manager loops or aborts when creating and deleting workspaces or when displaying the CDE Window List. • Fixes a memory leak problem in the Window Manager.
Patch 904.00 OSF510-487	**Patch:** Security (SSRT0713U) **State:** Supersedes patches OSF510-189 (395.00), OSF510-327 (638.00) This patch corrects the following: • 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. Compaq has corrected this potential vulnerability. • Fixes problems that may prevent a correct configuration table entry from being written to the binary error log on some systems, and may cause binlogd to display error messages on others. • 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.
Patch 909.00 OSF510-580	**Patch:** Fix to allow manual removal of Persistent Reserves **State:** New This patch allows the manual removal of Persistent Reserves from HSV110 if neccessary.
Patch 911.00 OSF510-576A	**Patch:** Security (SSRT0779) **State:** New A potential security vulnerability has been discovered where, under certain circumstances, SNMP services can stop functioning.
Patch 913.00 OSF510-576B	**Patch:** Security (SSRT0779) **State:** New A potential security vulnerability has been discovered where, under certain circumstances, SNMP services can stop functioning.
Patch 915.00 OSF510-581	**Patch:** Update for SCSI CAM Utility Program **State:** New This patch updates /sbin/scu, the SCSI CAM Utility Program. It adds support for Persistent Reserve for HSV110 as well as the display of 128-bit WWIDS.

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

Patch 917.00 OSF510-594	**Patch:** LAT setup does not handle inittab file as CDSL **State:** New. Supersedes patch OSF510-328 (544.00), 546.00 (OSF510-330) This patch corrects the following: • Fixes a problem in latsetup when the directory "/dev/lat" is not found. • Fixes a problem when latsetup does not handle the /etc/inittab file as a Context Dependent Symbolic Link (CDSL). • Fixes a problem where latsetup will fail if the system file /etc/inittab is not installed as a Context Dependent Symbol Link (CDSL).
Patch 919.00 OSF510-591	**Patch:** Savecore prematurely terminates crash dump recovery **State:** New This patch corrects a problem where savecore may prematurely terminate crash dump recovery on partitions larger than 4GB.

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

Patch 921.01 OSF510-568	**Patch:** Security (SSRT1-40U, SSRT1-41U, SSRT1-42U, SSRT1-45U) **State:** Supersedes patches OSF510-070 (132.00), OSF510-135 (330.00), OSF510-174 (250.00), OSF510-173 (252.00), OSF510-025 (59.00), OSF510-042 (156.00), OSF510-048 (53.00), OSF510-010 (60.00), OSF510-014 (62.00), OSF510-015 (151.00), OSF510-087 (152.00), OSF510-060 (154.00), OSF510-011 (11.00), OSF510-032 (22.00), OSF510-006 (23.00), OSF510-007 (24.00), OSF510-008 (25.00), OSF510-049 (26.00), OSF510-030 (27.00), OSF510-012 (28.00), OSF510-023 (29.00), OSF510-047 (30.00), OSF510-039A (32.00), OSF510-059 (86.00), OSF510-065 (93.00), OSF510-073 (94.00), OSF510-084 (95.00), OSF510-063 (96.00), OSF510-053 (97.00), OSF510-050 (98.00), OSF510-064 (99.00), OSF510-035 (100.00), OSF510-062 (101.00), OSF510-089 (103.00), OSF510-095 (163.00), OSF510-094 (165.00), OSF510-101 (167.00), OSF510-097 (176.00), OSF510-119 (177.00), OSF510-110 (178.00), OSF510-124 (179.00), OSF510-175 (180.00), OSF510-078 (181.00), OSF510-159 (182.00), OSF510-196 (183.00), OSF510-107A (184.00), OSF510-126 (185.00), OSF510-182 (186.00), OSF510-201 (187.00), OSF510-213 (188.00), OSF510-168 (189.00), OSF510-212A (190.00), OSF510-211 (191.00), OSF510-111 (192.00), OSF510-184 (193.00), OSF510-188 (194.00), OSF510-099 (195.00), OSF510-149 (196.00), OSF510-206A (197.00), OSF510-136 (198.00), OSF510-209 (199.00), OSF510-140 (200.00), OSF510-117 (201.00), OSF510-192 (202.00), OSF510-163 (203.00), OSF510-155 (204.00), OSF510-194 (205.00), OSF510-122 (206.00), OSF510-157 (207.00), OSF510-134 (208.00), OSF510-129 (209.00), OSF510-181 (210.00), OSF510-109A (211.00), OSF510-180 (212.00), OSF510-092 (213.00), OSF510-167 (214.00), OSF510-158A (215.00), OSF510-179 (216.00), OSF510-178 (217.00), OSF510-068 (218.00), OSF510-199 (219.00), OSF510-156 (220.00), OSF510-169 (221.00), OSF510-162 (222.00), OSF510-200 (224.00), OSF510-224 (399.00), OSF510-144 (328.00), OSF510-069 (92.00), OSF510-125 (366.00), OSF510-204 (409.00), OSF510-351 (410.00), OSF510-343 (411.00), OSF510-275 (412.00), OSF510-277 (413.00), OSF510-313 (414.00), OSF510-362 (415.00), OSF510-377 (416.00), OSF510-353 (417.00), OSF510-229 (418.00), OSF510-302 (419.00), OSF510-232 (420.00), OSF510-251 (421.00), OSF510-365 (422.00), OSF510-341 (423.00), OSF510-241 (424.00), OSF510-218 (425.00), OSF510-321 (426.00), OSF510-294 (427.00), OSF510-360 (428.00), OSF510-345 (429.00), OSF510-259 (430.00), OSF510-299 (431.00), OSF510-372 (432.00), OSF510-231 (433.00), OSF510-296 (434.00), OSF510-339 (435.00), OSF510-293 (436.00), OSF510-304 (437.00), OSF510-230A (438.00), OSF510-354 (439.00), OSF510-305 (440.00), OSF510-228 (441.00), OSF510-355 (442.00), OSF510-237 (443.00), OSF510-227 (444.00), OSF510-306 (445.00), OSF510-202 (446.00), OSF510-383 (447.00), OSF510-282 (448.00), OSF510-272 (449.00), OSF510-352 (450.00), OSF510-287 (451.00), OSF510-316 (452.00), OSF510-311 (453.00), OSF510-346 (454.00), OSF510-314 (455.00), OSF510-356 (456.00), OSF510-303 (457.00),

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

Patch 921.01 continued	OSF510-295 (458.00), OSF510-292 (459.00), OSF510-335 (460.00), OSF510-291 (461.00), OSF510-281 (462.00), OSF510-279 (463.00), OSF510-320 (464.00), OSF510-243 (465.00), OSF510-278 (466.00), OSF510-288 (467.00), OSF510-263 (468.00), OSF510-255 (469.00), OSF510-248 (470.00), OSF510-268 (471.00), OSF510-369 (472.00), OSF510-357 (473.00), OSF510-374 (474.00), OSF510-332 (475.00), OSF510-310 (476.00), OSF510-364 (477.00), OSF510-333 (478.00), OSF510-358 (479.00), OSF510-271 (480.00), OSF510-290 (481.00), OSF510-235 (482.00), OSF510-298 (483.00), OSF510-326 (484.00), OSF510-269 (485.00), OSF510-249 (486.00), OSF510-350 (487.00), OSF510-337 (488.00), OSF510-323 (489.00), OSF510-264 (490.00), OSF510-307 (492.00), OSF510-389 (639.00), OSF510-388 (640.00), OSF510-386 (641.00), OSF510-387 (642.00), OSF510-390 (644.00), OSF510-413 (645.00), OSF510-402 (647.00), OSF510-439 (650.00), OSF510-438 (651.00), OSF510-440 (653.00), OSF510-185 (331.00), OSF510-112 (333.00), OSF510-317A (511.00), OSF510-331 (634.00), OSF510-139 (374.00), OSF510-045 (45.00), OSF510-029 (47.00), OSF510-085 (109.00), OSF510-342 (579.00), OSF510-024 (49.00), OSF510-508 (656.00), OSF510-435 (657.00), OSF510-437 (658.00), OSF510-431 (659.00), OSF510-396 (660.00), OSF510-528 (661.00), OSF510-404 (662.00), OSF510-366 (663.00), OSF510-443 (664.00), OSF510-512A (665.00), OSF510-506 (666.00), OSF510-462 (667.00), OSF510-497 (668.00), OSF510-461 (669.00), OSF510-521 (670.00), OSF510-448 (672.00), OSF510-421 (673.00), OSF510-409 (674.00), OSF510-423 (675.00), OSF510-449 (676.00), OSF510-459 (677.00), OSF510-518 (678.00), OSF510-470 (679.00), OSF510-425 (680.00), OSF510-491 (681.00), OSF510-415 (682.00), OSF510-505 (683.00), OSF510-525 (684.00), OSF510-368 (685.00), OSF510-513 (686.00), OSF510-410 (687.00), OSF510-478 (688.00), OSF510-468 (689.00), OSF510-538 (690.00), OSF510-418 (691.00), OSF510-545 (692.00), OSF510-403 (693.00), OSF510-502 (694.00), OSF510-471 (695.00), OSF510-519 (696.00), OSF510-444 (697.00), OSF510-428 (698.00), OSF510-420 (699.00), OSF510-457 (700.00), OSF510-526 (701.00), OSF510-484 (702.00), OSF510-504 (703.00), OSF510-476 (704.00), OSF510-514 (705.00), OSF510-500 (706.00), OSF510-405 (707.00), OSF510-498 (708.00), OSF510-453 (709.00), OSF510-527 (710.00), OSF510-517 (711.00), OSF510-533 (712.00), OSF510-466 (713.00), OSF510-483 (714.00), OSF510-488 (715.00), OSF510-474 (716.00), OSF510-436 (717.00), OSF510-522 (718.00), OSF510-549 (719.00), OSF510-446 (720.00), OSF510-445 (721.00), OSF510-492 (722.00), OSF510-546 (723.00), OSF510-536 (724.00), OSF510-494 (725.00), OSF510-490 (726.00), OSF510-479 (727.00), OSF510-482 (728.00), OSF510-463 (729.00), OSF510-503 (730.00), OSF510-537 (731.00), OSF510-496 (732.00), OSF510-509 (733.00), OSF510-507 (734.00), OSF510-348 (735.00), OSF510-542 (736.00), OSF510-424 (737.00), OSF510-380 (738.00), OSF510-469 (739.00), OSF510-442 (740.00), OSF510-376 (742.00), OSF510-550 (906.00), OSF510-106 (308.00), OSF510-044 (55.00), OSF510-115 (266.00), OSF510-165 (268.00), OSF510-239 (526.00), OSF510-267 (528.00), OSF510-451 (788.00), OSF510-373 (789.00), OSF510-447 (790.00), OSF510-382 (791.00), OSF510-524 (792.00), OSF510-541 (908.00), OSF510-397 (794.00), OSF510-565 (905.00), OSF510-607 (907.00)

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

Patch 921.01 continued	This patch corrects the following problems: • Fixes an NFS file locking race. • Corrects the problem with write errors seen on soft-mounted NFS filesystems. The error received is: NFS3 RFS3_WRITE failed for server ncinfs: RPC: Server can't decode arguments • Corrects a problem where a race condition in NFS can result in a kernel memory fault. • Fixes a problem where threads can hang while renaming files on NFS mounted file systems. • This patch avoids tagged-file induced automount requests in AutoFS. • This patch is required in order to use the SuperDLT1 tape drive. • Fixes a problem encountered on a heavily loaded HSG80, in which a device may become unavailable to other cluster members if a cluster node crashes at the same time an error occurs on that device. • Prevents panics from occurring if AdvFS detects corruption in the per-fileset frags file and attempts to work around the corruption. • Fixes AdvFS memory mapped file support so that it honors the noatimes and readonly mount options when updating file timestamps. • A kernel memory fault can occur on an SMP machine when one thread is extending a clone frags file and another thread does a stat system call on a file with a frag. • Provides an improvement to AdvFS performance when the first bytes of user data (and subsequent storage requests) is written to a domain. • Corrects read-ahead behavior for AdvFS for both local and NFS reads. Read performance is increased by approximately 10% with the addition of this patch. This patch does not include any correctness fixes. • Fixes a problem on AlphaServer GS80, GS160, and GS320 systems where, under a specific set of unlikely circumstances, it is possible for Revision 4 PCA hardware to falsely report PCI hung bus errors. This will cause a uncorrectable hardware machine check and operating system panic. This patch must be installed if the hardware configuration includes any Revision 4 PCA (IOP to PCI bus) adapters. • Fixes a kernel memory fault which can occur during scheduler load balancing on a NUMA system.

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

Patch 921.01 continued	• Fixes a panic that occurs in madvise() when called with MADV_DONTNEED when running in lockmode 4. • Improves performance of HPTC programs on GS-series NUMA machines. • Fixes a kernel memory fault which can happen when all the physical memory is in use. • Fixes a problem seen in a cluster when one member whose boot partition is on a device whose SCSI wwid changes while the node is down. • Corrects a failure that is seen as a user_cmd timeout. • Fixes a kernel memory fault when accessing a shared text segment after or during load balancing on a NUMA system. • Fixes a bug that, when fuser -k is issued on a dismounted NFS mount point in which a process is running, a hang will occur. • 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. Compaq has corrected this potential vulnerability. • Improves webserver performance, fixes an IPV6 related crash, and a hang in soclose(). • Fixes problems for threaded applications running on Tru64 UNIX V5.1 • sbin/dd has been made non-threaded. This is to avoid problems while installing patches that are incompatible with the running kernel. /usr/bin/dd is not affected by this patch. • Fixes SPECweb99 httpd hangs in umc_get_page() routine waiting for the page. • Includes performance fixes for systems doing raw I/O, raw async I/O, and systems with large disk farms (high disk count). • Enables the getconf command to return the abbreviated vendor name correctly. • Provides the device driver for a new graphics card. • Fixes a problem where some network-based multimedia applications will cause a kernel memory fault when exiting. • Provides support for the DEGPA-TA (1000BaseT) Gigabit Ethernet device. • Fixes a potential deadlock on systems using shared memory segments and granularity hints. This can occur when allocating a gh region larger than the available free memory. • Improves UDP performance by removing an unneeded lock from the UDP output path. • Fixes a panic in in_pcbfree() when NFS is implemented over TCP. • Fixes a lock contention for multiple writers which would use 100% of CPU time. This problem has been seen when running Oracle database doing Table Creates. • Resolves hang-like behavior when LSM volumes are used to create AdvFS domain volumes. The default preferred I/O byte transfer size may be too large and needs to be set lower. • Fixes periodic slowdowns seen on large systems that are consuming large amounts of memory due to file I/O. These changes make the reclaiming of memory in use for file buffers more efficient. There is also a fix for a lock timeout seen on the vdIoLock because of a large number of buffers on the SmoothSync queues.

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

Patch 921.01 continued	• Fixes a race condition which could result with either a Kernel Memory Fault or a Kernel Unaligned Access in one of the AdvFS I/O queue manipulation routines. • Fixes inaccuracy problems when using setrlimit/getrlimit with a threaded application. • Addresses multiple issues for the KZPCC family of RAID Array 2000 (RA2000) controllers: – Errors seen when concurrent opens are issued to separate logical partitions on the same logical device. – Change to the preferred chunk size from 16 KB to 64 KB which may increase data transfer rates. • Fixes a hang seen while running collect and the vdump utility. This patch prevents the hang in tok_wait from occurring. • Prevents stat(), lstat(), fstat(), statfs(), fstatfs(), getmntinfo(), and getfsstat() from returning EOVERFLOW errors for programs compiled on Tru64 UNIX V4.0 or earlier. • Fixes a problem where threads can hang in x_load_in-mem_xtnt_map(). • Fixes a kernel memory fault when writing to /proc, while anon_rss_enforce is set to 2. • Fixes an issue with lightweight wiring of pages and shared memory regions. • Fixes a system panic when the system has at least one AdvFS domain and the system is configured for lockmode=4 kernel lock statistics collection. • Corrects some I/O rate fluctuations and thread unresponsiveness that had been seen when vm free pages dropped to a low level and used pages were being recycled. • In laboratory testing Compaq has observed that, under certain circumstances, a possibility exists that inconsistent data may be written to disk on some Tru64 UNIX V5.0A and V5.1 systems running AdvFS and direct I/O. Compaq became aware of this possibility only during laboratory testing. To our knowledge, no customer has experienced this problem. Compaq is alerting customers to this potential problem as a precautionary measure. The conditions under which this potential problem may occur are as follows: – An application writes to a file using AdvFS direct I/O and the file had previously been opened for normal I/O (which by default is cached). – Some but not all of the pages are still resident in Unified Buffer Cache (UBC) memory. Invalid data could occur when a single direct I/O write spans multiple AdvFS pages, and some, but not all, of the pages are still in the UBC. If the file has been opened only for direct I/O and remains open for direct I/O, the problem does not exist. Applications that use direct I/O, such as Oracle, could be affected. • This patch addresses two types of system crashes: – Crash caused by VM hash corruption, kernel memory fault. – Crash caused by a lock hierarchy violation.

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

Patch 921.01 continued	• Fixes a problem with the driver for Gigabit Ethernet adapters (DEGPA-FA and DEGPA-TA) which prevented its use in a NetRAIN (Redundant Array of Independent Network Adapters) set. • Fixes a problem where the setgid bit of a directory was not being set when created if its parent directory has the setgid bit set. • Fixes issues with memory allocation attributes. • Fixes a bug in the POSIX Threads Library for Tru64 UNIX V5.1 where a terminating thread did not properly clear an enabled floating point unit, causing an invalid floating point state on the next thread that was run. • Fixes several virtual memory algorithms related to the allocation and freeing of pages within the kernel. • Fixes panics which can occur if a signal is sent to a multi-threaded task in which one or more threads are calling exit() or exec(). • Fixes the corruption of the CAM hardware database when using hwmgr. This typically can result in a kernel memory fault when the database is being written to disk after a hwmgr operation. • Corrects an AdvFS panic which can occur during a rmfset operation, causing the following panic string: rbf_delete_int: can't find bf attributes • Fixes an issue with some remote ioctls for tape/changer drivers not working in a cluster. • Fixes a panic which comes from a page fault on a user buffer while already holding the write lock. • Fixes a bug in the POSIX Threads Library for Tru64 UNIX V5.1 that would result in a DECthreads error return of EINVAL from the pthread mutex API routines. This error would be seen only when the thread stack had been user defined/changed, specifically seen when using the user level context switching (ucontext) routines. • Fixes a problem in which the system panicked with a kernel memory fault while the class scheduler was being configured. • Fixes cluster hangs where I/O stops, and a hwmgr -view -clu command does not return. However, the systems will respond to pings. This is caused by the ubc_memory_purge in routine cfs_putpage being blocked when doing FSOP_PUTPAGE.

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

Patch 921.01 continued	• Fixes the following system panics: – A "simple_lock: lock already owned by cpu" panic when anon_rss_enforce is non-zero and lockmode is set to 4. This remove occurs when a process, whose RSS (resident set size; the number of pages a process can have in memory) limit is exceeded tries to expand its heap. – A "panic: vm_page_activate: already active" panic that can occur on a system during memory shortages. – An "mcs_lock: no queue entries available" panic that can occur on a GS160 system. This is caused by an abandoned page mistakenly being reclaimed off the the 0/O hash. The page is then removed off a UBC free list where two stale page pointers were connected, hereby connecting the ACTIVE and INACTIVE list. When attempting to deactivate pages (move them from the ACTIVE queue to the INACTIVE queue) an INACTIVE page is encountered, which causes an inadvertant failure to unlock the page. Continued attempts to deactivate INACTIVE pages results in the lock queue being filled. This can also cause a "kernel memory fault" panic. • Fixes a problem in which a heavy load placed on an HSG80 can disable the device. • Fixes a timing window where flushing data to disk can be incomplete when a system is going down, or if more than one thread calls reboot() without first going through shutdown, /sbin/reboot, or /sbin/halt. • Fixes a system crash that could occur when calling nmadvise. • 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. Compaq has corrected this potential vulnerability. • Eliminates a kernel memory fault in AdvFS. • Fixes multiple problems with SCSI tape handling including improvements to backup procedures, SCSI passthrough, an increase to the local IO size for transfers, a fix for a system crash that can occur during a bus reset, and a fix for a panic with the following panic string: PWS_CCB_QUE_REMOVE: CCB NOT ON ANY LIST • Fixes a system hang caused by netisr queue corruption due to a race condition that is primarily encountered by third party drivers and layered products that call schednetisr_nospl(). • Fixes a lockmode 4 panic in netisr_del_rad where netisr_del_rad attempted to release a lock it did not hold. • Corrects the use of Granularity Hints in a threaded application program. • Fixes a problem with writing out crashdumps on systems with their swap on FibreChannel.

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

Patch 921.01 continued	• Fixes a kernel memory fault and invalid memory ifetch panic which can occur in AlphaServer SC systems running Quadrics' RMS software. • Fixes a bug in the POSIX Threads Library for Tru64 UNIX V5.1 that would result in a DECthreads Bugcheck and process termination. Threaded applications might encounter this problem when pthread_kill() is used on a thread that is marked as blocked in the kernel. • Corrects the behavior of the FIONBIO, FIOASYNC, and FIONREAD ioctls in a cluster environment. These commands would fail, returning ENOTTY when they should have succeeded. • Fixes a problem in which the system call fcntl(fd, F_DUPFD, 15) fails with "too many files" even after fd limits have been increased. • Corrects two problems with the scheduler: – Enables NUMA load balancing in other processor sets, then the default processor set (pset 0). – Enables the processor to do load balancing for multi-threaded applications. • Provides support for activating temporary data logging on a mount point. • Fixes a hang in the ufs filesystem. • Fixes kernel build failures due to an undefined ss_sched function. • Fixes a problem with the execution of interpreter programs failing with a "file not found" error if the the total space used by the environment variables and command arguments is close to a multiple of 8K (page size). • Provides full KZPCC support in Version 2.0 of the i2o block storage driver. Restriction: For TCR-V5.1 installations, KZPCC support is restricted to data-only service; devices on the KZPCC controller cannot be used for system or boot partitions in a cluster. Additionally, this patch fixes the problem where extraneous console messages will appear when hardware is added or deleted. • Fixes a performance problem with V5.1 where threads doing large I/O transfers could spend excess time in ubc_page_alloc(). • Fixes nmadvise with a modification to VM to allow migration of shared memory. • Provides functionality to support EMC storage boxes that support Persistent Reserves (SCSI command set) as defined by the final SCSI specification. • Fixes a kernel memory fault in GS series systems which have mixed revision PCI adapters. • Fixes the following two issues: – "u_anon_free: page_busy" system panic when using System V shared memory locked by a single process. – Failures ranging from uninitialized simple_lock panics, kernel memory fault panics, and process hangs on GS320/160/80 systems configured with at least one memory less quad. • Fixes the automount handling of the "nogrpid" option. • Fixes a network problem where a system can hang during a route command.

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

Patch 921.01 continued	• Addresses two problems with the ee driver for DE60x Ethernet cards. These problems affect all Tru64 UNIX systems containing ee cards: – Fixes a race condition where the card could stop receiving packets from the network under rare circumstances. – Fixes the lan_config user options -x and -s. • Fixes some problems seen with loading and unloading dynamic drivers. • Fixes a couple of problems in NFS that can cause a kernel memory fault during NFS server shutdown. • Corrects a problem with ICMP redirect processing which resulted in incorrect ICMP redirect messages. • Fixes a kernel memory fault when performing asynchronous input/output over sockets. • 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 problem with sendmsg and rcvmsg that prevented 9i/RAC from being able to use UDP as its transport. With this patch, correct operation of sendmsg and rcvmsg is restored when dealing with atomic protocols by not truncating send but to treat as a 32 bit length. • Fixes a kernel memory fault in mount -o extend. • Provides a script, /usr/sbin/clone_versw_undo, that will allow a user to remove the directio cloning 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. • Fixes a rare panic in the driver for the DE600/DE602 10/100 Ethernet adapter. • Fixes data inconsistency problems that can be seen on clusters that are NFS clients.

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

Patch 921.01 continued	• Fixes a misconfiguration of vm_free_target at the boot time when this parameter is added to /etc/sysconfigtab. • Fixes problems seen with the loading and unloading of dynamic drivers. • Fixes a kernel memory fault in tcp_rad_slowtimo. This patch also fixes a kernel memory fault in soclose() before calling soabort for listener sockets. • 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 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. • Prevents the error message "local HSM Error: msgsvc: socket close failed" from being generated when an application closes the socket with return state 0. • Fixes a problem in which activity to a disk that is connected to an HSG80 will hang if the disk is removed and reinserted. • Prevents a potential hang due to external NFS servers. • Fixes a panic in ubc_page_release while running direct I/O. The fix ensures that even pre-allocated pages get flushed, thus preventing an lru corruption. • Fixes a problem where, when using VX1 graphics module, the mouse cursor disappears when moved along the left and topmost edges. • Fixes a system panic with "malloc_check_checksum: memory pool corrution" message. • Corrects several problems in kernel routing: – Fixes a panic when deleting an IP address. – Fixes a panic when performing IP re-configuration. – Adds interface route on address configuration. • Corrects a problem in the virtual file system that could cause panic with the panic string "kernel memory fault."

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

Patch 921.01 continued	• Fixes a bug between mcs_unlock and mcs_lock_try on the same CPU, causing the mcs_unlock to hang. • Ensures that if an AdvFS file is opened for both O_DIRECTIO and O_APPEND, threads racing to append data to the file will be correctly synchronized, and all data will be appended to the file. • Fixes a bug in virtual memory that can cause a kernel memory fault. • Fixes a condition where the smoothsync thread, in attempting to flush dirty buffers for memory-mapped files, would also flush buffers for non-memory-mapped files. This did not cause any errors, but could cause more I/O than necessary. • 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 kernel panic with the following message: bs_invalidate_rsvd_access_struct: bad access struct • Makes the balance and rmvol programs in AdvFS more interruptible by supplying a new option (-i). It also avoids wasting extent map entries and avoids a kmf in overlay_xtnt_map. • Fixes the following problems: – The system may hang while attempting to replace a component that is used in a redundant configuration. – The system may experience a kernel memory fault when an I/O path is removed. Just before the panic occurs, you may see: Jun 24 16:21:05 tstsys vmunix: DDR - Warning: Device has no "name" - for Jun 24 16:21:05 tstsys vmunix: Vendor ID : Product ID: • Fixes a kernel memory fault when using open a command hwmgr -delete component -id 3. • Fixes a problem that would cause a process to hang because the process was unable to exit. • Eliminates superfluous AutoFS auto-mount attempts during rolling upgrade. These attempted auto-mounts slow down certain operations and leave the AutoFS namespace polluted with directories prefixed with .old.. • Fixes some problems with the mkdir -p command when executed on automount directories. • Fixes a problem where a long-running kernel thread in AdvFS could cause a cluster timeout and subsequent panic. It also fixes a simple_lock timeout panic.

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

Patch 921.01 continued	• Corrects a problem with the network code which resulted in some tcp packets having the wrong checksums. This could result in dropped connections. • Fixes lock time issues, UBC performance problems, and provides AdvFS and UFS performance improvemnts in platforms, other than Wildfire, with low memory. • Fixes a problem with AdvFS that, when mounting the filesystem with option -o dual a panic is caused. • 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". Compaq has corrected this potential vulnerability. • Fixes a panic seen on a cluster that displays the panic string "mcs_lock: time limit exceeded". In the dump you will see both dyn_lock_bucket and dyn_hash_obtain_chain. • Fixes a kernel memory fault caused by AutoFS. • The table() system will not abort connections properly if a tcb hash table number is greater than 1. • Corrects an "mcs_lock: time limit exceeded" panic when moving processors to/from processor_sets. • Fixes a bug that can cause performance problems for certain applications when the sysconfigtab parameter ipc:sem_broadcast_wakeup is set to 0. • Fixes several directIO problems seen when using the aio interface. The symptoms include a kernel memory fault, and an aio condition that causes a live_dump to be generated. • 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: – 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 inconsistancies caused by fault path for large shared memory regions prematurely releasing a hold on a page it just locked. This can cause variety of problems including user program errors and system panics. – A "simple_lock: time limit exceeded" system panic that occurs if very large (8MB or larger) System V Shared memory regions are in use.

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

Patch 921.01 continued	• Corrects the problem of a simple lock timeout due to POSIX timers and also corrects some inaccuracies of the POSIX realtime timers. • Fxes a kernel memory fault in msg_rpc_trap. • Fixes a problem where the I/O transfer rate can suddenly drop when writing to a hole in a large file in an AdvFS domain, when a volume in that domain becomes full. • Fixes a panic with the following error message: panic: cfsdb_assert • Prevents a hang in msfs_cfs_flush_and_invalidate() when running defragment on a cluster. • Fixes a problem such that applications that directly manipulate memory buffer pointers get correct results. This problem is exhibited when using Tarantella Enterprise 3 application server software to run applications. The UNIX system will hang, requiring a power shutdown and system reboot to recover. • This patch will fix panics generated by whole-file flushes of metadata files. Symptoms include: CLUSTER BOOT PANIC: SIMPLE_LOCK: UNINITIALIZED LOCK KMF IN ADVFS_PAGE_BUSY() DURING RECOVERY PROCESSING PANIC WHEN MOUNTING ADVFS FILE SYSTEM ADVFS CLUSTER ROOT DOMAIN GOT CORRUPTED • Replaces the system panics caused by "Can't clear bit twice" with a domain panic. • The mkfdmn command now works with the -V3 and -p options. This prevents a core dump from being generated. This is a rare situation that was seen by code inspection. • Domain panics that were inadvertently removed from bs_frag_alloc() have been replaced. • A potential security vulnerability has been discovered in the kernel where, under certain circumstances, a race condition can occur that could allow a non-root user to modify any file and possibly gain root access. • Fixes a problem in which netisr_add() can erroneously return an EEXIST error. This problem can manifest as "Framework error: connection problems" messages from X.25 applications.

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

Patch 921.01 continued	• Addresses a panic situation in IN_PCBREF and a change to tcp_deletetcb to prevent a crash. • Corrects several CAM errors including: passthru IOCTL fails with EIO (CAM_BUSY) problem; RESERVATION CONFLICT driver BUSY problem; enforce super user only access for SCSI passthru. • Fixes a cluster problem where opening a file after open/close of its clone deadlocks the AdvFS thread. • Adds unified wait support in conjunction with clustered RDG multichannel wait flag fix to allow for more efficient processing by Oracle processes. • Fixes a problem where network interfaces can appear unresponsive to network traffic. • Corrects a CFS problem where the data on an AdvFS clone fileset may get overwritten as an unexpected side effect of using directio. The problem occurs when the program issuing the directio open is running on a CFS client AND the fileset involved has been cloned AND a rewrite occurs involving pages not yet modified since the creation of the clone. • Fixes mbuf memory corruption when using ICS/TCP. • Fixes a problem with vm_faults against anon objects mapped by multiple map entries. • Corrects the problem of a thread deadlocking against itself under the following conditions: – 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 performance problem and the results are large performance increases in configurations where more than 8 tapes are supported on a FibreChannel (usually behind an MDR or FCTCII).

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

Patch 921.01 continued	• Fixes a problem in kernel threads where multi-threaded applications were allowed to start running prior to virtual memory mapping swapin. This was prevented by adding a flag to mark when the map is swapped out and no thread swapins can occur until this flag is cleared. • Fixes a problem in the Virtual Memory subsystem where a process hangs and cannot be killed. This problem only happens on NUMA systems. • Contains fixes that ensure FibreChannel system configurations can properly identify boot and swap devices required to obtain crash dumps. This patch requires that FibreChannel systems which utilize FibreChannel devices for boot and swap be properly configured. • Fixes a panic of "malloc_leak: free with wrong type" when using kmem-debug-protect. • Fixes an issue where Sybase reports "Error: 1613" and "host process disconnected" errors. • A threaded section of application code can crash when using granularity hints (GH). • Ensures that certain invariants within the kernel concerning clone maps are maintained. It maintains consistency and correctness of the clone maps. • Fixes a problem that can cause a "kernel memory fault" panic in load_from_shadow_rec(). • Fixes incorrect usage of UNMOUNT_TRY_READ in AutoFS. • Fixes a bug that can cause a panic when a system is powering down. • Fixes the following problems using the hwmgr command: KMF FTX_DONE_URDR: BAD FTX UNALIGNED KERNEL SPACE ACCESS FROM KERNEL MODE KMF FROM HWC_LOOKUP_DEVT_SAFE HWCC_JACKET_RTN: BAD CALL TO KCH HWCC_EVAL_REQUEST: INFALLIBLE PROPOSAL RETURNED ERROR HWCC_JACKET_RTN: INFALLIBLE PROPOSAL RETURNED ERROR • Prevents lock hierarchy violations due to putpage/migrate interaction.

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

Patch 921.01 continued	• Fixes a problem where an AdvFS direct I/O read can cause a "kernel memory fault" system panic. The problem occurs when the following two conditions are met: – One of the pages cannot be read. – The I/O request is not an even multiple of 512 bytes. • Allows POSIX semaphores/msg queues to operate properly on a CFS client. • 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. • Fixes a system panic with the panic string: "lock_terminate: lock held". This is being caused by the table call which, when accessing an open file table from another task, was not doing the proper locking. • A potential security vulnerability has been discovered in networking where, under certain circumstances, a remote system can take over packets destined for another host. • Fixes a problem where the UBC subsystem fails to purge pages because of bound purge_thread. • Fixes the following system panics: Kernel Memory Fault in function sth_close_fifo() when closing a vnode that belongs to a FIFO simple_lock: time limit exceeded in spec_reclaim • Fixes a problem in which a TCP socket can continue to receive data with no application running. • Corrects a problem where the network subsystem sometimes sends a null TCP packet when a connection is reset. • 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 panic with "xfer_hole_stg: unaligned kernel access" or "xfer_hole_stg: kernel memory fault" messages.

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

Patch 921.01 continued	• Fixes an "RDG unwire panic" when running with RDG and GH chunks. • Adds support for future version of Emulex FibreChannel adapter. • Fixes the following tape drive problems: – Tape devices in multi-path configurations unexpectedly rewind or go offline. (Multi-path 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 offline 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 • Fixes errors generated by syscheck when NFS is not configured. • Upgrades sys_check to V120. • 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). Compaq has corrected this potential vulnerability. • Provides the /usr/lbin/mkstemp program which allows the mechanism to create a secure temporary file. • Corrects a hang that can be seen on multi-CPU systems using NFS-over-TCP. The SMP race is seen between the nfs_tcp_input and the nfs_tcp_thread functions. • Adjusts the sleep time for slower robot tape changers to allow them time to replace a tape. • Fixes a kernel panic caused by btcreate when it generated scripts to recreate LSM volumes on restore operations. • Fixes a problem where the device special files are not being created by btextract.

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

Patch 921.01 continued	• Fixes a problem with the 400ms delay upon network cable reinsertion which could lead to temporarily held drivers. • Fixes the processing of export lists with a / (slash) in them. • Fixes a situation where a failed open to a device will cause an error that the device cannot be deleted using hwmgr. • 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 to an AdvFS file using asynchronous I/O will hang if it is completed on a memoryless QBB. • Fixes a kernel memory fault due to a bug in kernel code. • Corrects the problem where attempts to delete psets can hang the system. • Fixes a "u_shm_oop_deallocate: reference count mismatch" due to a bug in locking mechanism when gh_chunks are in use. • Corrects problems with USB causing panics under heavily stressed systems. • Corrects kernel memory inconsistencies against the 4096-byte bucket when SWCC is running and a control port is deleted. • Fixes a timing window that caused queue inconsistencies. • Corrects an issue with mmap()ed files on a NFS mounted filesystem. Changes to an mmap()ed file were not being immediately seen. • Installs DECthreads V3.18-138 which fixes problems that may affect threaded programs running on Tru64 UNIX V5.1. • 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. • Fixes numerous problems of accessing de-allocated and freed vnodes. • Fixes a problem which can result in a panic, hang, or corruption from vnode deallocation during an unmount. • Fixes the following problems: – Prevents HSG80 controller crashes. – Fixes cam_logger error message problems during cluster boot. – Fixes DRD problems and persistent reservation problems. • Fixes AdvFS synchonization problems with lingering I/O messages and domain deactivation or rmvol. It also fixes problems caused by certain kmem_debug settings and AdvFS handling of freed memory. • 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. • Processes triggering stack growth with anon_rss_enforce set to 2, and exceeding the set resident memory limit hang or panic. • During filesystem 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.

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

Patch 921.01 continued	• Fixes several problems with AutoFS: – A problem resulting in a panic in Clusters. – Intercept point not created due to busy mounton directory. – Inadvertent unmounts of locally mounted file systems. – An intercept point rendered unusable after an error during an auto-mount attempt. – Eliminates error messages concerning property lists seen via certain utilities such as vdump. – AutoFS auto-mounts will now occur when utilities name intercept points defined via indirect map entries. • Fixes a problem where opens would fail when running under heavy IO load with the KZPCC. • Corrects a problem whereby clocks on systems could move backwards after subsequent relocations of the root file system using cfsmgr. • Prevents a panic in bs_derefpg. • Fixes a problem where the tape changer is only accessible from member that us the DRD server for the changer. • Fixes locking on retry case for multi-threaded select/poll. A panic with the following panic string is indicative of this problem: PANIC: "thread_block: simple lock owned" • Corrects a problem relating to the negative lookup cache behavior that causes a negative lookup result to hide the results of a successful create operation. • 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. Compaq has corrected this potential vulnerability. • Prevents the error message “No such file or directory” from autofsd when the asterisk character (*) is used. • Addresses a problem where file locks set on a file from an NFS client may not be properly released at the NFS server. This could cause any future lock requests (local lock requests, as well as lock requests from NFS clients) for that particular file to block indefinitely. • Corrects an NFS hang when the delayed option is used with the mount command. • Eliminates AdvFS domain panics for filesystems served remotely on a local disk, when the server node is shut down. • 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. Compaq has 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 improper file access. Compaq has corrected this potential vulnerability. • Fixes locking problems in vclean(). • Prevents a possible extent map corruption when multiple volumes are full. • Prevents a race in msfs_umount.

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

Patch 921.01 continued	• Corrects a CFS problem that could cause a panic with the panic string of "CFS_INFS full". • Fixes a problem where socket-based applications can hang in soclose(). • Corrects a problem with the counters maintained for the NetRAIN virtual interface. • Fixes a system panic and/or data inconsistencies caused by changing fifo parameter pipe-databuf-size while fifo operations are in flight. • Fixes a problem where multi-threaded processes may hang in timed condition waits (pthread_cond_timedwait()) when running realtime system contention scope threads. • Fixes a panic experienced while task swapping • Fixes a "kernel memory fault" panic on NUMA systems because of corrupt UBC LRU. • Fixes a problem that causes bugchecks from applications running DECthreads. • Fixes a problem with poor interactive performance including hanging commands and logins, and random drops in I/O rates when writing many large files. • Fixes a kernel memory fault on a UFS filesystem from calc_extentmap. • Fixes and enhances Tru64 UNIX to support Encore realtime software. • Fixes a problem where I/O writes may not update attributes properly. • Fixes the CEH bus/target and lun number when lun > 127. • Corrects a kernel memory fault on multiple CPU systems when two or more CPUs find an AdvFS problem at the same time. • 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 defragment, balance, migrate, rmvol, or mssh. • Fixes a potential "kernel memory fault" panic in the Virtual Memory subsystem on SMP systems. • Fixes a crash in hwc space when lockmode is equal 4 and add support to get devt information from user space. • Contains AlphaServer ECC Enhancements for DTAG error logging • Fixes reservation conflicts in cdisk_rec_tur_done. • A potential data inconisistency problem has been discovered 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. A second possible symptom is incomplete flushing of user data when an fsync() is issued or an O_[D]SYNC write is performed. Compaq has corrected this problem. • A second problem has been discovered in which a call to fsync() or fsyncdata() may return to an application before all of the data is safely on disk. Compaq has corrected this problem.

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

Patch 921.01 continued	• Fixes a problem where decreasing the smoothsync_age does not always have an effect. • 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. Compaq has corrected this potential vulnerability. • Fixes a problem that causes a system panic when a program calls sendfile(2) to access a file via NFS. • Fixes a problem where tape and changer devices on fibre could occassionally return an incorrect offline status. • Fixes a problem where, when running ssh v2.4.0 and v2.4.1, users will see a problem executing ls in sftp and when uploading public key using ssh-pubkeymgr. • Adds support to AutoFS for auto-mount map syntax for replicated servers in Sun's style. It also fixes a problem wherein mount options were not being applied to replicated servers after the first one listed. • Corrects problems where NFS can deadlock and also corrects an AdvFS problem where EIOs are returned by AdvFS to NFS. • Fixes the object selection audit style (auditmask -s obj_sel) so that files removed from, or added to, a directory which is flagged (auditmask -x dir_path) to be monitored for access or modification generates an audit event. • Corrects a problem where multi-volume AdvFS V3 domains exhibit I/O errors (not attributable to hardware). The same problem also causes a failed mkfset due to ENO_XTNTS. • Corrects a race condition which could result in a failure to set the modification time of a file. This occurs only on a UFS filesystem. • Fixes a problem where, after a system crash, on reboot there is a domain panic with the following stack trace: 7 domain_panic 8 ftx_bfmeta_rec_redo 9 ftx_recovery_pass 10 ftx_bfdmn_recovery 11 bs_bfdmn_activate 12 bs_bfdmn_tbl_activate 13 bs_get_dmntbl_params 14 msfs_real_syscall 15 msfs_syscall 16 syscall • Addresses a kernel memory fault panic in malloc_thread(). • 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 a crash that occurs when disk controllers are restarted repeatedly. • Fixes a potential problem where system responsiveness may be affected.

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

Patch 921.01 continued	• Addresses two problems with the alt driver for DEGPA Gigabit Ethernet adapters. These problems affect all Tru64 systems using al" with vMAC or NetRAIN: – Fixes vMAC support. Prior to this patch, vMAC has not worked with DEGPA. – Prevents two DEGPA adapters from getting the same MAC address in a NetRAIN configuration. • Two problems are corrected for non-NUMA systems: – 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 a race condition in the class scheduler that could cause a Kernel Memory Fault. • Addresses problems with the NFS portmap and mountd daemons. These problems are cluster-specific, and could result in services that register with portmap becoming unusable. • Contains several fixes to the disk driver: – Corrects a panic due to an I/O barrier failure. – Corrects memory inconsistencies due to the use of a path structure that is deleted before being used. – Corrects a problem where path lists could become unstable if driver recovery was in progress. – Corrects a panic due to a lock hierarchy ordering problem. • Addresses a data inconsistency that can occur when a CFS client reads a file using direct I/O that was recently written to. • Corrects the problem where the DLI queue stalls when there is no traffic in the TCP/IP or HDLC stacks. • Fixes a problem where previously, storage allocation for a file opened for directIO 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. • Fixes the predictable TCP Sequence Number. • Fixes a potential CFS deadlock. • Fixes an incorrect priority return value from sched_getparam().

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

Patch 921.01 continued	• Fixes a problem with device descriptor references in clusters. The halting of one cluster node would cause the entire cluster to crash • Prevents an AdvFS metadata inconsistency in the event of a system crash. • Fixes C shell processing problems in the new zh_CN.GB18030 locale. • 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. Compaq has corrected this potential vulnerability. • Fixes a possible handling problem with multibyte character boundary conditions in ksh script processing. • Fixes two ksh problems that occur in multi-byte Asian locales. • Fixes a problem in which /usr/bin/ksh hangs for certain scripts that contain wait(1). • fixes a problem with ksh. When a ksh menu is started from within a user's .profile, ksh will not stop when the Telnet session is stopped. • Fixes a problem with the Cshell (csh) so that it now correctly recognizes the backslash (\) meta character. • Corrects a problem in which ksh fails to substitute the tilde (~) character for a user's home directory after an assignment using the # or % characters has been used. • Corrects two problems with csh(1): – If a non-root user performed an ls(1) with wild card 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 nonomatch 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. • 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. Compaq has corrected this potential vulnerability. • The following changes were made: – Shell inline input files are more secure. – sh noclobber and new constructs are added. – The mkdir system call is updated. • Fixes an Asian language processing problem under the Korn shell. • Fixes a problem with multi-threaded applications that can cause the application to consume 100% of the CPU usage time. • This patch is required to support the undo of the version-switched cloning patch when installed via the nonroll patch capablity.

3

Summary of TruCluster Software Patches

This chapter summarizes the TruCluster software patches included in Patch Kit-0005.

Table 3–1 lists patches that have been updated.

Table 3–2 provides a summary of patches in Patch Kit-0005.

Table 3–1: Updated TruCluster Software Patches

Patch IDs	Change Summary
Patches 119.00, 122.00, 127.00, 129.00	New
Patches 5.00, 7.00, 32.00, 70.00, 72.00, 17.00, 42.00, 43.00, 45.00, 83.00	Superseded by Patch 85.00
Patches 28.00, 41.00, 46.00, 47.00, 49.00, 86.00, 87.00, 88.00	Superseded by Patch 90.00
Patch 120.00	Superseded by Patch 122.00
Patches 11.00, 30.00, 67.00, 69.00	Superseded by Patch 131.00
Patches 78.00, 80.00, 141.00	Superseded by Patch 143.00
Patches 66.00, 123.00, 125.00	Superseded by Patch 148.00
Patches 15.00, 33.00, 34.00, 35.00, 36.00, 37.00, 39.00, 73.00, 74.00, 75.00, 77.00, 132.00, 133.00, 134.00, 135.00, 136.00, 137.00, 138.00, 140.00	Superseded by Patch 150.00
Patches 2.00, 13.00, 18.00, 19.00, 20.00, 21.00, 22.00, 23.00, 24.00, 26.00, 50.00, 51.00, 52.00, 53.00, 54.00, 55.00, 56.00, 57.00, 58.00, 59.00, 60.00, 61.00, 62.00, 64.00, 82.00, 91.00, 92.00, 93.00, 94.00, 95.00, 96.00, 97.00, 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, 115.00, 117.00, 146.00	Superseded by Patch 152.01

Table 3–2: Summary of TruCluster Patches

Patch IDs	Abstract
Patch 4.00 TCR510DX-001	**Patch:** Fix for Cluster Alias Manager system management tool **State:** Existing This patch fixes the Cluster Alias Manager system management tool from crashing and displaying errors.
Patch 9.00 TCR510-001	**Patch:** Initializing the MC-API results in system crash **State:** Existing This patch fixes a problem where on the AlphaServer GS160 systems, initializing the MC-API results in the system crashing with a "kernel memory fault" message.

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

Patch 85.00 TCR510-107	**Patch:** Fixes memory hang **State:** Supersedes patches TCR510-002 (5.00), TCR510-003 (7.00), TCR510-023 (32.00), TCR510-042 (70.00), TCR510-039 (72.00), TCR510-018 (17.00), TCR510-028 (42.00), TCR510-052 (43.00), TCR510-043 (45.00), TCR510-095 (83.00) This patch corrects the following: • Fixes an occasional cluster hang which can occur after a Memory Channel error. • Fixes a kernel memory fault which occurs in the ics_mct_ring_recv() routine. The kernel memory fault is seen when a node is booting into the cluster, and can occur on the booting node or on another node. • Fixes a problem in ICS where ring_recv() does not properly handle a change in channel numbers. The fix will, in turn, improve validation of the connection structure on node joins. • Fixes the way communication errors occur on clusters such that a down node will not declare all other nodes dead. • Fixes the problem that causes a panic with error message "CNX QDISK: Yielding to foreign owner with quorum" caused by a long running thread, ICS/MCT receive thread, which defers other kernel threads from accessing the CPU. • Eliminates unnecessary rail failovers in vhub configurations and removes rmerror_int diagnostic messages. • Fixes an issue which causes all cluster nodes to hang or panic if a Wildfire is halted via the halt button. • Fixes a panic that is caused in a clustered environment that has the following error message: rm_request_on_bad_prail • Prevents an "ics_mct: Error from establish_RM_notification_channel" panic on clusters. • Fixes four problem situations: – When a physical MC rail goes offline. – When the master failover node goes offline during a failover. – How ICS handles the resend situation when MC errors take place. – Failing over due to parity errors increasing beyond the limit. • Fixes hangs and increases performance of memory channel ICS operation.

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

Patch 90.00 TCR510-087	**Patch:** Fixes a panic in clua_cnx_unregister **State:** Supersedes patches TCR510-019 (28.00), TCR510-029 (41.00), TCR510-041 (46.00), TCR510-048 (47.00), TCR510-037 (49.00), TCR510-091 (86.00), TCR510-082 (87.00), TCR510-066 (88.00) This patch corrects the following: • Fixes the cluamgr command where it will display the alias status even if no cluster member has joined the alias. • Fixes a problem in which RPC requests to the cluster alias may fail with "RPC timeout" message. • Fixes a cluster node hang from in_pcbnotify. • Fixes a problem that a rebooted node not able of sending messages to the cluster alias. • Fixes multiple networking issues within a cluster environment: – 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.memebrX. – 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 a problem when the cluster alias subsystem does not send a reply to a client that pings a cluster alias address with a packet size of less than 28 bytes. • Fixes a memory corruption panic which could occur after a member joins the cluster or after adding a new cluster alias to one or more of the members. • Fixes a problem with cluster alias selection priority when adding a member to an alias. • Fixes a panic in clua_cnx_unregister where a TP structure could not be allocated for a new TCP connection.
Patch 119.00 TCR510DX-002	**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. Compaq has corrected this potential vulnerability.
Patch 122.00 TCR510-063	**Patch:** cfsmgr works correctly with upper case member names **State:** New. Supersedes patch TCR510-070 (120.00) This patch corrects the following: • Correct a cfsmgr error "Not enough space" when attempting to relocate a file system with a large amount of disks. • Allows cfsmgr to work correctly with upper and mixed case member names.

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

Patch 127.00 TCR510-092	**Patch:** Using a cluster as a RIS server causes a panic **State:** New This patch addresses two problems: • A panic caused by a known problem, using a cluster as a RIS server. • A fix to RIS/DMS serving in a TruCluster.
Patch 129.00 TCR510-071	**Patch:** EVM cluster-wide event may cause a panic **State:** New This patch fixes a problem that, under very heavy loads in a cluster, could cause the system to panic when duplicating a cluster EVM event.
Patch 131.00 TCR510-104	**Patch:** Fix for Oracle 9i hang **State:** Supersedes patches TCR510-007 (11.00), TCR510-024 (30.00), TCR510-036 (67.00), TCR510-049 (69.00) This patch corrects the following: • Corrects a problem in which the RDG subsystem will stop sending messages even though there are messages which are deliverable. • Fixes an incorrect display of the following warning message at boot time: rdg: failed to start context rcvq scan thread • Fixes a kernel memory fault with the RDG autowiring mechanism, also seen as a “pte not valid” crash. • Adds a multichannel wait flag to pid_unblock. • Contains performance enhancements. • Fixes a problem with RDG whereby broadcast packets can interact with the context receive queue. • Closes a timing window that can cause Oracle 9i to hang when a remote node in the cluster goes down.
Patch 143.00 TCR510-085	**Patch:** Panic in distributed lock mgr deadlock detection code **State:** Supersedes patches TCR510-033 (78.00), TCR510-047 (80.00), TCR510-061 (141.00) This patch corrects the following: • Fixes an Oracle process hang if a node fails after receiving a “rsbinfo”message. • Fixes a DLM problem where two processes could take out the same lock. • Fixes a panic in dlm when another node in the cluster is halted. • Fixes a panic in the distributed lock managed deadlock detection code.

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

Patch 148.00 TCR510-121	**Patch:** CAA applications not failing over **State:** Supersedes patches TCR510-027 (66.00), TCR510-067 (123.00), TCR510-110 (125.00) This patch corrects the following: • For systems running TruCluster Server V5.1 with the following configurations: – Tapes and/or media changer devices used as CAA resources. – A combination of tapes, media changers, and network interfaces used as CAA resources. • Fixes a problem that prevents CAA from updating the state of any of the above resources when connectivity to the corresponding device (tape, media changer, or network) is lost or restored. • Fixes a situation when CAA daemon on a clustered system crashes and dumps core. • Fixes the major problems of CAA applications not failing over during a node shutdown and caad hang condition at startup. • Corrects the inability to start and stop CAA resources. When started they will go to the unknown state and never start. The problem is nondeterministic. Several CAA resources may be started before the problem is seen.

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

Patch 150.00 TCR510-115	**Patch:** Failover does not occur properly **State:** Supersedes patches TCR510-005 (15.00), TCR510-021 (33.00), TCR510-009 (34.00), TCR510-016 (35.00), TCR510-011 (36.00), TCR510-022 (37.00), TCR510-012 (39.00), TCR510-035 (73.00), TCR510-038 (74.00), TCR510-030 (75.00), TCR510-034 (77.00), TCR510-109 (132.00), TCR510-108 (133.00), TCR510-094 (134.00), TCR510-065 (135.00), TCR510-084 (136.00), TCR510-105 (137.00), TCR510-106 (138.00), TCR510-090) (140.00) This patch corrects the following: • Fixes two TruCluster problems: – If a Quorum disk is manually added by the command clu_quorum -d add, the disk becomes inaccessible because the PR flag is not being cleaned up. The same command will work in the next reboot. – A cluster member cannot boot under a specific hardware setup. The CFS mount fails because of the PR flag is not cleaned up. • Addresses the need for IOCTL for remote DRD, adds clean up for failed remote closes for non-disks, fixes error returns on failed tape/changer closes, and fixes tape deadlock experienced in netbackups. • Fixes an issue with a tape/changer failing to correctly report a close failure of a device in a cluster environment. • Fixes a problem which results in a system panic while doing tape failovers. • Fixes a node panic during fiber port disables. • Fixes an issue with a tape/changer giving back “busy on open” if a close from a remote node failed. • Provides the TCR portion of the functionality to support EMC storage boxes that support Persistent Reserves (SCSI command set) as defined by the final SCSI specification. • Fixes an issue with requests being stuck on a failed disk in a cluster. • 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 an issue with the hwmgr -delete command that causes a panic in a cluster. • Fixes the KZPCC controller problem seen when deleting a Virtual Drive using SWCC and adding the same drive back can result in the disk being unaccessible. • Fixes several problems with the device request dispatcher (drd) kernel subsystem, including cluster hangs, kernel memory faults, reboot problems, node recovery problems, and device failover problems. • Fixes cluster hangs and panics due to I/O problems. • Fixes a problem where the tape changer is only accessible from member that's the drd server for the changer. • Fixes a race condition problem when multiple unbarrierable disks failed at the same time. • Fixes a problem where CAA applications using tape/changers as required resources will not come ONLINE (as seen by caa_stat).

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

Patch 150.00 continued	• Fixes a kernel memory fault in drd_open. • Fixes the following problems: – Prevents HSG80 controller crashes. – Fixes cam_logger error message problems during cluster boot. – Fixes DRD problems and persistent reservation problems. – Fixes problems with drdmgr not responding to a failover disk. – Fixes a domain panic in a cluster when a file system is mounted on a disk accessed remotely over the cluster interconnect.

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

Patch 152.01 TCR510-123	**Patch:** Security (SSRT0691U) **State:** Supersedes patches TCR510-004 (2.00), TCR510-006 (13.00), TCR510-026 (18.00), TCR510-020 (19.00), TCR510-013 (20.00), TCR510-015 (21.00), TCR510-017 (22.00), TCR510-014 (23.00), TCR510-025 (24.00), TCR510-008 (26.00), TCR510-056 (50.00), TCR510-050 (51.00), TCR510-054 (52.00), TCR510-057 (53.00), TCR510-046 (54.00), TCR510-040 (55.00), TCR510-031 (56.00), TCR510-032 (57.00), TCR510-051 (58.00), TCR510-060 (59.00), TCR510-044 (60.00), TCR510-053 (61.00), TCR510-045 (62.00), TCR510-058 (64.00), TCR510-064 (82.00), TCR510-077 (91.00), TCR510-100 (92.00), TCR510-098 (93.00), TCR510-081 (94.00), TCR510-072 (95.00), TCR510-073 (96.00), TCR510-075 (97.00), TCR510-083 (98.00), TCR510-093 (99.00), TCR510-096 (100.00), TCR510-069 (101.00), TCR510-088 (102.00), TCR510-076 (103.00), TCR510-079 (104.00), TCR510-086 (105.00), TCR510-089 (106.00), TCR510-078 (107.00), TCR510-099 (108.00), TCR510-097 (109.00), TCR510-102 (110.00), TCR510-101 (111.00), TCR510-103 (112.00), TCR510-074 (113.00), TCR510-080 (114.00), TCR510-062 (115.00), TCR510-068 (117.00), TCR510-127 (144.00), TCR510-123 (146.00) This patch corrects the following: • 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. Compaq has corrected this potential vulnerability. • Provides a small TPC-C performance optimization to cfsspec_read for reporting TPC-C single node cluster numbers. • When attempting to roll a patch kit on a single member cluster without this patch, the following error messages will be seen when running the postinstall stage: *** Error*** Members '2' is NOT at the new base software version. *** Error*** Members '2' is NOT at the new TruCluster software version. • During backup stage of clu_upgrade setup 1, clu_upgrade is unable to determine the name of the kernel configuration file. • clu_upgrade does not check the availabilty of space in /, /usr, and /usr/i18n. • During the preinstalled phase, clu_upgrade will ignore a no answer when the user is prompted, during an error condition, whether they wish to continue. • clu_upgrade incorrectly assumes that if the directory /usr/i18n exists, then it is in its own file system. • After the clu_upgrade clean phase, the final step of clu_upgrade, no message is displayed that leads the user to believe they have completed the upgrade. Only the prompt is returned and the clu _upgrade -completed clean command reports that the clean had not completed. • clu_upgrade can display "Could not get property..." and "...does not exist" type of error messages during the undo install phase. • The clu_upgrade undo switch command, after completing a clu_upgrade switch command, should display an error message instead of claiming it has succeeded. • Fixes a problem with disaster recovery whereby the node being restored will hang on boot.

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

Patch 152.01 continued	• Corrects a problem in which a cluster may panic with a "cfsdb_assert" message when restoring files from backup while simultaneously relocating the CFS server for that file system. • Corrects a problem in which a cluster member can panic with the panic string "cfsdb_assert" when a NFS V3 TCP client attempts to create a socket using mknod(2). • Corrects a problem in which a cluster member will panic with the patch string "lock_terminate: lock held" from cinactive(). • Fixes a hang seen while running collect and the vdump utility. This patch prevents the hang in tok_wait from occurring. This also prevents a cfsdb_assert panic that contains the following message: Assert Failed: (tcbp->tcb_flags & TOK_GIVEBACK) == 0 • Prevents a cfsdb_assert panic from occurring in the cfs block reserve code. The system is most likely running process accounting that will receive this type of panic. • Provides performance enhancements for copying large files (files smaller than the total size of client's physical memory) between a CFS client and server within the cluster. • Corrects a token hang situation by comparing against the correct revision mode. • Fixes a bug in the cluster filesytem that can cause a kernel memory fault. • Eliminates superfluous AutoFS auto-mount attempts during rolling upgrade. These attempted auto-mounts slow down certain operations and leave the AutoFS namespace polluted with directories prefexed with ".Old..". • Fixes memory leak in cfscall_ioctl(). • Fixes a panic with the following error message: panic: cfsdb_assert • Contains corrections required for proper operation of Oracle 9i with Tru64 UNIX/TruCluster 5.1. The problems corrected include: – Processes hanging when using Cluster File System/Direct I/O feature. – Improper handling of direct I/O to an AdvFS fileset if a clone fileset was already in use, potentially resulting in an inconsistent backup. – Using ls -l, the Cluster File System file attribute could be seen inconsistently from the server and client members. For example, a file's mode could be seen differently from the server and the client. – A file opened for Direct I/O on the Cluster File System server may inappropriately be opened in non-direct I/O mode by a client. – Oracle processes hanging due to shutting down one cluster member. – A problem with the Cluster File System which could cause a cluster system to panic with the panic string "kernel memory fault" in the routine mc_bcopy(). – A problem with Cluster File System which could cause a cluster member to panic with the panic string "uiomove: mode." This problem could cause Oracle multi-instance data bases to crash with the message similar to the following: ORA-27063: skgfospo: number of bytes read/written is incorrect

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

Patch 152.01 continued	• Fixes data inconsistency problems that can be seen on clusters that are NFS clients. • Frevents a cfsdb_assert panic from occurring in cfs_reclaim. This panic has been seen while running ensight7. • Prevents a potential hang due to external NFS servers. • Provides a warning to users installing 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. • Prevents a potential hang that can occur on a CFS failover. • Allows POSIX semaphores/msg queues to operate properly on a CFS client. • Allows the command cfsstat -i to execute properly. • Corrects a problem which can cause cluster members to hang, waiting for the update daemon to flush /var/adm/pacct. • Fixes a potential CFS hang on defragment. • Fixes a possible “Kernel Memory Fault” panic on racing mount update/unmount/remount operations for the same mount point. • Fixes a possible “Kernel Memory Fault” in function ckidtokgs. • Fixes possible "cfs_add_mount() - database entry present" panic and possible multinode reboot hang which shows the following message: WARNING: RETRYING TO LOCK THE BOOT PARTITION DEVICE • Fixes two race conditions in Cluster Mount support: – One results in a transient mount failure. – The second might result in a kernel memory fault panic during mount. • Fixes two AutoFS problems: – AutoFS is unable to establish an intercept point when mounton directory is busy. – Fixes an “Unaligned Kernel Access” panic in cfs_vget_fhp(). • Fixes a panic that would occur during the mount of a cluster file system on top of a non-cluster file system. • Prevents a “Kernel Memory Fault” panic during unmount in a Cluster or during a planned relocation. • Corrects a "cfsdb_assert" panic which can occur following the failure of a cluster node.

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

Patch 152.01 continued	• Addresses three CFS problems: – A kernel memory fault in the CFS read-ahead code. – A deadlock in the CFS read-ahead code. – A potential data inconsistency problem which could occur when a filesystem becomes 100% full. • Enforces the rule that mounting on a server-only file system makes the new mount server-only. • Fixes two race conditions: – Between cluster root failover and mount which results in a kernel memory fault. – Between failover-related cleanup and bootup-time mount processing, which results in deadlock and hangs the new node. • Eliminates a Kernel Memory Fault panic during node shutdown. • Addresses a problem in CFS where, under certain conditions, CFS would temporarily change the value of p_pid of the current running process. The result of this could break certain pid-based hashing algorithms in the kernel, as well as advery affect certain kernel debugging tools. • Fixes a race condition during cluster mount which results in a transient ENODEV seen by a name space lookup. • Addresses a problem where a file's attributes (owner, group, mode, etc) could become inconsistent cluster-wide. • Fixes a PANIC: CFS_ADD_MOUNT() - DATABASE ENTRY PRESENT panic when a node re-joins the cluster. • Addresses a problem where CFS may not properly invalidate cached access rights when a change is made to a file's property list. • Fixes a race condition between node shutdown and unmount, and ensures that all file sets from an AdvFS domain mounted as server_only get unmounted when the server node is shut down. • This patch addresses two cluster problems: – Hung unmounts, possibly seen as hung node shutdowns. – A cfsdb_assert panic in cfs_tokmsg(). • Fixes the assertion failure ERROR != ECFS_TRYAGAIN. • Corrects a CFS problem that could cause a panic with the panic string of "CFS_INFS full". • Fixes several potential CFS panics. • Fixes functional problems dealing with CFS direct I/O and CFS block reservation. • Fixes a possible panic on boot if mount request is received from another node too early in the boot process. • 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(). • This patch prevents the following panic: cms_kgs_callback_thr: in use already set on non-initiator

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

Patch 152.01 continued	• Fixes a potential CFS deadlock. • 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. • Fixes a hang during cluster unmount which results in the blocking of all further mounts and unmounts. • 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.