Tru64 UNIX and TruCluster Server Version 5.1B

Patch Summary and Release Notes for Patch Kit 2

April 2003

This manual describes the contents of Patch Kit 2 for the 5.1B version of the Tru64 UNIX operating system and TruCluster Server Software products. It provides special instructions for installing individual patches.

See the *Technical Updates for Tru64 UNIX Patch Kits* for information about restrictions and problems that may have been discovered since the release of this kit. See the *Patch Kit Installation Instructions* for information about installing or removing patches, baselining, and general patch management.

Hewlett-Packard Company
Palo Alto, California

Contents

About This Manual

1 Tru64 UNIX Patches

1.1 Release Notes 1–1
1.1.1 Required Storage Space 1–1
1.1.2 Problems Uninstalling the Patch Kit 1–1
1.1.3 Required Action for Removing Patch 16.00 1–2
1.1.4 ES47/ES80/GS1280 Time Loss 1–2
1.1.5 Enabling envmond on AlphaServer ES47/ES80/GS1280 systems (Patch 452) 1–2
1.1.6 New Security Feature (Patch 872.00) 1–3
1.1.7 Possible Problem when Processing Many Command Parameters 1–4
1.1.8 Loading Firmware from a BOOTP Server 1–4
1.1.9 Enhancements to pmgrd Daemon and collect Utility (Patches 779.00 and 781.00) 1–5
1.1.9.1 Performance Manager Metrics Server Daemon (pmgrd) 1–5
1.1.9.2 collect Utility 1–6
1.1.10 File System Management Applications Enhanced (Patch 825.00) 1–6
1.1.11 Changes to tar, pax, and cpio Behavior (Patch 643.00) 1–6
1.1.12 Changes to vdump and vrestore Allow Larger Record Sizes (Patch 685.00) 1–7
1.1.13 Problem Seen on Systems with Smart Array Controller 1–7
1.1.14 Broken Links Reported During Baselining 1–7
1.1.15 Russian Keyboard (Patch 339.00) 1–8
1.1.16 Panics May Occur on Multi-CPU Systems 1–8
1.1.17 General and Problem Information for AlphaServer ES47, ES80, and GS1280 Systems 1–8
1.1.17.1 CPU Offline Restrictions 1–8
1.1.17.2 Problem with Capacity-on-Demand Process 1–9
1.1.17.3 Hardware SCSI Bus Errors 1–9
1.1.17.4 Repeated Reboots May Cause Panic 1–9
1.1.18 Caution on Updating to Version 5.1B with DEGXA NICs 1–9
1.1.19 Tuning the NFS Server Duplicate Request Cache (Patch 838.00) 1–9
1.2 Summary of Base Operating System Patches 1–10

2 TruCluster Server Patches

2.1 Release Notes 2–1
2.1.1 Required Storage Space 2–1
2.1.2 AlphaServer ES47 or AlphaServer GS1280 Hangs When Added to Cluster 2–1
2.1.3 No-Roll Procedure Cannot Be Used to Remove Patch Kit 2–2
2.1.4 Updates for Rolling Upgrade Procedures 2–2
2.1.4.1 Noncritical Errors 2–2
2.1.4.2 Procedure for Simultaneous Upgrades 2–2
2.1.4.3 Unrecoverable Failure Procedure 2–3

2.1.4.4 Do Not Add or Delete OSF, TCR, IOS, or OSH Subsets During Roll 2–3
2.1.4.5 Undo Stages in Correct Order 2–3
2.1.4.6 Ignore Message About Missing ladebug.cat File 2–3
2.1.4.7 clu_upgrade undo of Install Stage Can Result in Incorrect File Permissions 2–3
2.1.4.8 Missing Entry Messages Can Be Ignored During Rolling Patch 2–4
2.1.4.9 Relocating AutoFS During a Rolling Upgrade on a Cluster .. 2–4
2.1.5 Additional Steps Required When Installing Patches Before Cluster Creation 2–5
2.1.6 When Taking a Cluster Member to Single-User Mode, First Halt the Member 2–5
2.1.7 Problems with clu_upgrade switch Stage 2–6
2.2 Summary of TruCluster Software Patches 2–6

About This Manual

This manual contains information specific to Patch Kit 2 of the Tru64 UNIX operating system and TruCluster Server software products for Version 5.1B. It briefly describes the patches contained in this kit and provides information you should be aware of when installing certain patches.

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 Provides information about the Tru64 UNIX patches included in this kit.

Chapter 2 Provides information about the TruCluster Server software patches included in this kit.

Related Documentation

In addition to this manual, the following documentation may be helpful in the patching process:

- *Technical Updates for Tru64 UNIX Patch Kits*

  This document reports any information about restrictions and problems that may have been discovered since the release of this and other patch kits.
- Tru64 UNIX and TruCluster Server *Patch Kit Installation Instructions*
- *Patching Best Practice*
- The `dupatch`(8) reference page, which describes the use of `dupatch` from the command line. This reference page is installed when you install the `dupatch` tools.
- Tru64 UNIX *Installation Guide*
- Tru64 UNIX *System Administration*
- TruCluster Server *Cluster Installation*
- TruCluster Server *Cluster Administration*
- Release-specific installation documentation

Patch Process Resources

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

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

  http://ftp1.support.compaq.com/public/unix/
- To view or print the latest version of the *Patch Kit Installation Instructions* or the *Patch Summary and Release Notes* for a specific patch kit:

  http://h30097.www3.hp.com/docs/patch/index.html
- To visit our main support page:

http://h71025.www7.hp.com/support/home/index.asp

- To visit the Tru64 UNIX homepage:

  http://h30097.www3.hp.com/

Reader's Comments

We welcome any comments and suggestions you have on this and other Tru64 UNIX manuals.

You can send your comments in the following ways:

- Fax: 603-884-0120 Attn: HCTO Information Development, 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:

  Hewlett-Packard Company
  HCTO Information Development Manager
  ZK03-3/Y32
  110 Spit Brook Road
  Nashua, NH 03062-9987

Please include the following information along with your comments:

- The full title of this document.
- The section numbers and page numbers of the information on which you are commenting.
- The version of Tru64 UNIX that you are using.
- If known, the type of processor that is running the Tru64 UNIX software.

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

1

Tru64 UNIX Patches

This chapter provides information about the patches included in Patch Kit 2 for the base operating system. It also includes any general information about working with these patches.

This chapter is organized as follows:

- Section 1.1 provides release notes that are specific to the Tru64 UNIX patches in this kit, as well as release notes that are of general interest.
- Section 1.2 provides brief descriptions of the Tru64 UNIX patches included in this kit.

1.1 Release Notes

This section provides release notes that are specific to the Tru64 UNIX patches in this kit, as well as release notes that are of general interest.

1.1.1 Required Storage Space

Approximately 250 MB of temporary storage space is required to untar the base and TruCluster components of this patch kit. We recommend that this kit not be placed in the `/`, `/usr`, or `/var` file systems because doing so may unduly constrain the available storage space for the patching activity.

The following permanent storage space is required to install the base component of this patch kit:

- Approximately 76.9 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.
- Approximately 78.2 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.
- Approximately 1142 KB of storage space is required in `/var/adm/patch/doc` for patch abstract and README documentation.
- Approximately 78247 KB of storage space is needed in `/usr/sbin/dupatch` for the patch management utility.

See Section 2.1.1 for information on space needed for the TruCluster Server patches.

1.1.2 Problems Uninstalling the Patch Kit

If you made the following changes to your system after installing the patch kit, you will have to undo those changes before you can uninstall the patch kit:

- If you changed your hardware configuration (for example, by adding a new disk), the system configuration that existed prior to installing the patch kit might not recognize the new devices or may not provide the necessary support for them.
- If you added new cluster members, the new members will not have an older state to revert to if you attempt to uninstall the patch kit.

To uninstall the patch kit, do the following:

1. Remove all new hardware and new cluster members that you added after installing the patch kit.
2. Run `dupatch` to uninstall the patch kit.
3. Verify that the patch kit was successfully uninstalled.

You can now add the cluster members you removed and reinstall the hardware you removed, as long as the support for it existed in the pre-patched system. You can also reinstall the patch kit.

1.1.3 Required Action for Removing Patch 16.00

If you delete Patch 16.00, you must run the `/etc/dn_fix_dat.sh` script prior to rebooting your system. If you are removing the patch from a cluster, you must run the script on each cluster member before you reboot the member. See Section 2.1.3 for more information.

1.1.4 ES47/ES80/GS1280 Time Loss

The ES47, ES80, and GS1280 AlphaServers may experience a time loss as a result of console callbacks for environmental information if the server's firmware is lower than V6.4-12.

Updating your firmware to V6.4-12 or higher will keep the problem from occurring or correct the problem if it has occurred.

If your firmware is lower than V6.4-12, the problem is experienced if one or both of the following conditions exists:

- The system manager uses the following `hwmgr` utility commands:

```
# hwmgr -view devices
# hwmgr -view hierarchy
```

- The Environmental Monitoring daemon, `envmond`, is running.

Workarounds to the problem are as follows:

- Do not use the `hwmgr -view` commands
- Modify either one of the following files:
  - `/etc/rc.config`

    Turn off environmental monitoring by changing the entry `ENVMON_CONFIGURED=1 to ENVMON_CONFIGURED=0`

    You can also use the `envconfig` utility to modify the `/etc/rc.config` file. See `envconfig`(8) for information on using the utility.

  - `/etc/sysconfigtab`

    Go to the end of the file and add the following line and add the following line:

```
marvel_srvmgmt: MV_Env_Support = 0
```

    You must remove this setting after you install firmware V6.4-11 or higher.

You must reboot your system for these settings to take effect.

1.1.5 Enabling envmond on AlphaServer ES47/ES80/GS1280 systems (Patch 452)

After installing Patch 452, the `envmond` daemon will be disabled on ES47, ES80, and GS1280 AlphaServers..

To enable environmental monitoring, change the entry `ENVMON_CONFIGURED=0 to ENVMON_CONFIGURED=1` in the `/etc/rc.config` file. You can do this using one of the following commands:

```
/usr/sbin/envconfig -c ENVMON_CONFIGURED=1
/usr/sbin/rcmgr set ENVMON_CONFIGURED 1
```

If your server's firmware is lower than V6.4-12, see Section 1.1.4 before enabling environmental monitoring.

1.1.6 New Security Feature (Patch 872.00)

Patch 872.00 provides a new security feature to prevent the execution of instructions that reside in heap or other data areas of process memory. The result is additional protection against buffer overflow exploits. This feature is similar in concept to Tru64 UNIX executable stack protection.

The new feature is implemented as a dynamic sysconfig tunable, `executable_data`, in the proc subsystem. The supported settings allow system administrators to cause requests from privileged processes for writable and executable memory to fail, or to be treated as a request for writable memory, and to optionally generate a message when such a request occurs.

In a buffer overflow exploitation, an attacker feeds a privileged program an unexpectedly large volume of carefully constructed data through inputs such as command line arguments and environment variables. If the program is not coded defensively, the attacker can overwrite areas of memory adjacent to the buffer.

Depending upon the location of the buffer (stack, heap, data area), the attacker can deceive these programs into executing malicious code that takes advantage of the program's privileges or alter a security-sensitive program variable to redirect program flow.

With some expertise, such an attack can be used to gain root access to the system.

Enabling the `executable_data` tunable changes a potential system compromise into, at worst, a denial-of-service attack. A vulnerable program may still contain a buffer overflow, but an exploit that writes an instruction stream into the buffer and attempts to transfer control to those instructions will fail, because memory protection will prohibit instruction execution from that area of memory.

Many applications never execute from the memory even though they unnecessarily request write-execute memory directly or as a result of an underlying function acting on their behalf. By substituting writable memory for the requested write-execute memory, the `executable_data` tunable allows such applications to benefit from the additional protection without requiring application modification. See `sys_attrs_proc`(5) for more information.

Before enabling `executable_data` (changing it from the default value of 0), you must run the `/usr/sbin/javaexecutedata` script. Otherwise, privileged java applications will fail in unpredictable ways. See `javaexecutedata`(8) for more information.

Note

The Java language interprets bytecode at runtime. Unless marked as exempt, privileged applications written in Java will receive an error when they attempt to execute instructions residing in the unexecutable memory. The manner in which these errors are handled is application-specific and thus unpredictable. This is why you

must run the /usr/sbin/javaexecutedata before you enable executable_data.

The following example demonstrates the failing behavior to expect for a privileged processes if execute_data is set to 53 but runs the /usr/sbin/javaexecutedata script. Other Java applications run with privilege may exhibit different (but still failing) behavior.

```
# java -classic -jar SwingSet2.jar
Process 1185 Invalid write/execute mmap call denied.
Process 1185 Invalid write/execute mmap call denied.
Process 1185 Invalid write/execute mmap call denied.
(...)
Process 1185 Invalid write/execute mmap call denied.
Process 1185 Invalid write/execute mmap call denied.
**Out of memory, exiting**
```

The following example demonstrates the failing behavior to expect for a privileged processes if execute_data is set to 37 but runs the /usr/sbin/javaexecutedata script. Other java applications run with privilege may exhibit different (but still failing) behavior.

```
# java -classic -jar SwingSet2.jar
Process 1185 Invalid write/execute mmap call modified.
Process 1185 Invalid write/execute mmap call modified.
 (...)
Process 1185 Invalid write/execute mmap call modified.
Process 1185 Invalid write/execute mmap call modified.
Process 1185 Invalid write/execute mmap call modified.
SIGSEGV   11*  segmentation violation
(...)
Abort (core dumped)
```

Certain privileged Pascal programs may also fail when executable_data is enabled. Such programs should also be marked as exempt, using the new chatr utility, included in Patch 872.00 and described as follows:

```
$chatr +ed enable priv_pascal_executable
  current values:
     64-bit COFF executable
     execute from data: disabled
  new values:
     64-bit COFF executable
     execute from data: enabled
```

See chatr(1)

1.1.7 Possible Problem when Processing Many Command Parameters

When running commands or scripts that must process a large amount of command parameters, your system may hang or you may see an error similar to this: /sbin/ls: arg list too long.

If this occurs, try rerunning the command or script after entering the following command to relax the command-line limits:

```
# sysconfig -r proc exec_disable_arg_limit=1
```

This kernel setting should not be used as a default. It should only be enabled when encountering a problem where the exec() argument size limit has been approached.

1.1.8 Loading Firmware from a BOOTP Server

The fwupgrade command has been modified in Patch Kit 2 to allow the specified firmware update image to be loaded from a BOOTP server in a connected network.

This process must use the bootpd daemon. The subset where the bootpd ships is optional, so OSFOBSOLETE540 must be installed.

Create a symbolic link from the shipping location of bootpd to the expected location:

```
# ln -s /usr/opt/obsolete/usr/sbin/bootpd /usr/sbin/bootp
```

You must manually create the bootptab file on the server. The following is an example of how to set up the bootptab file on the server for this procedure:

```
 # Example bootptab file for BOOTP support

.default1:\
:hd=/install/firmware:\
:sm=255.255.255.0\
:gw=16.69.255.1:

#
tab:tc=.default1:ht=ethernet:ha=08002b86f234:ip=16.69.222.42:
bobafett:tc=.default1:ht=ethernet:ha=0008c73a5a47:ip=16.69.222.48:
#
```

In this example, the directory /install/firmware was created on the bootp server. This directory must contain the firmware of the systems to be updated. The file names must match the entry on the fwupgrade command line. The firmware files must have read permissions, that is, 444.

You must edit the inetd.conf file so that the file name passed by fwupgrade is found by the console firmware. Edit the line /etc/inetd.conf file on the bootp server to look like following:

```
tftp   dgram   udp  wait    root /usr/sbin/tftpd tftp -r /install/firmware
```

Enable bootpd to start by removing the comment symbol (#) from the beginning of the line in the /etc/inetd.conf file;

```
bootps dgram   udp  wait    root /usr/sbin/bootpd  bootpd
```

See the fwupgrade(8), bootptab(4), and bootpd(8) reference pages for more information.

1.1.9 Enhancements to pmgrd Daemon and collect Utility (Patches 779.00 and 781.00)

The installation of Patch 779.00 and Patch 781.00 provides enhancements to the performance manager metrics server daemon, pmgrd, and the collect utility.

1.1.9.1 Performance Manager Metrics Server Daemon (pmgrd)

The following features have been added to pmgrd:

- Support for monitoring the disk I/O rates.

  Enables pmgrd to provide details on disk I/O rates, such as the average number of bytes transferred per second and the average number of transfers completed per second over the past 1 minute, 5 minutes, 30 minutes, and 60 minutes.

- Support for monitoring the AdvFS statistics.

  Enables pmgrd to provide the following types of details on AdvFS file systems:

  – The domain name
  – The fileset name
  – Number of files and blocks
  – Soft and hard limits of files
  – Soft and hard limits of blocks
  – The status of user and group quotas

– Grace time and fileset clone information

It can also provide AdvFS volume details such as available blocks, percentage of volume used, I/O consolidation mode, and the number of read/write blocks. The new MIB file `pmAdvfs.mib` has been added to provide these statistics.

The `collect` utility displays these new AdvFS statistics. (See Section 1.1.9.2).

As a result of the improvements made to `pmgrd`, we recommend that you use the SysMan Menu to manage AdvFS file systems rather than the `dtadvfs` graphical user interface and the `advfsd` daemon. To use SysMan Menu, select Storage -> File System Management Utilities -> Advanced File System (AdvFS) Utilities. You can also enter the following command:

```
# sysman advfs
```

See `pmgrd`(8) for more details. Patch 876.00 installs the revised `pmgrd`(8) reference page.

1.1.9.2 collect Utility

The following features have been added to the `collect` utility, which is updated from Version 2.0.0 to 2.0.5:

- AdvFS monitoring capability. (See Section 1.1.9 for a list of AdvFS metrics that are monitored.)

  Enables `collect` to report AdvFS volume I/O queue and fileset vnode operation statistics. You can specify the domain or fileset to be monitored, using the `-A` option.

- Viewing CPU and memory metrics on a per Resource Affinity Domain (RAD) basis.

  When run on a NUMA platform, enables `collect` to automatically retrieve CPU and memory metrics for each RAD.

See `collect`(8) for more details. Patch 876.00 installs the revised `pmgrd`(8) reference page.

1.1.10 File System Management Applications Enhanced (Patch 825.00)

Patch 825.00 enhances SysMan Menu file system management applications to significantly improve their performance.

1.1.11 Changes to tar, pax, and cpio Behavior (Patch 643.00)

When extracting or listing an archive using the `tar`, `pax`, or `cpio` commands, specifying a slash (/) at the end of argument will cause the command to act upon the directory and not the contents in the directory. For example:

```
# tar xvf filename.tar dir1/
```

When creating an archive with these commands, specifying multiple slashes will result in the placement of one slash for any directory entry in the archive header. Previously, specifying multiple slashes would put these slashes in the archive header. For example:

```
# tar cvf filename.tar dir1//////////
```

Specifying a single slash when creating the archive will cause `tar`, `pax`, or `cpio` to pick up all of the directory's contents. For example:

```
# tar cvf filename.tar dir1/
```

1.1.12 Changes to vdump and vrestore Allow Larger Record Sizes (Patch 685.00)

The `vdump` and `vrestore` programs have been tuned to work with higher record sizes up to 2048 KB. This provides a performance gain when doing backups of AdvFS domains, because the program will make larger tape records for the save sets.

Currently the maximum for the `-b` option is 64 1024-byte blocks. With this patch, the byte-block size is changed to 2048. By default, the `-b` option of 60 blocks per record remains unchanged..

The `vrestore` program still has the capability to autosize `vdump` archives.

1.1.13 Problem Seen on Systems with Smart Array Controller

This section describes the steps you should take if your system is configured with a Smart Array controller and you see the following event logged:

```
Host name: unx104
SCSI CAM ERROR PACKET
SCSI device class: CISS (Smart Array)
Bus Number: 6
Target Number: 4
Lon Number: 0
…
...
Event Information: Command timed out...resetting controller
```

If this occurs, take the following steps:

1. Create a file named `ciss.temp` with the following lines:

   ```
   ciss:
   ciss_throttle_threshold=5
   ```

2. Execute the following command:

   ```
   # sysconfigdb -m -f ciss.temp
   ```

3. Reboot your system:

   ```
   # shutdown -r now
   ```

1.1.14 Broken Links Reported During Baselining

When performing a baseline analysis with the `dupatch` utility, you will encounter the following message during Phase 4:

```
Phase 4 - Report changed system files and missing files
=======================================================

    This phase provides information to help you make choices later in
    this process.  It reports both 'missing' and files whose origin
    cannot be determined.  Some of these files may affect patch
    installation.  You will want to consider this information when you
    later make decisions in phase 5.

     * list of changed files with unknown origin:
      ----------------------------------------

  ./etc/lprsetup.dat                                          OSFPRINT540       UNKNOWN
  ./usr/share/doclib/annex/man/man3/Thread.3                  OSFTCLBASE540     UNKNOWN
      BROKEN HARDLINK TO ./usr/share/doclib/annex/man/man3/Tcl_ConditionNotify.3
  ./usr/share/doclib/annex/man/man3/Tcl_ConditionNotify.3     OSFTCLBASE540     UNKNOWN
      BROKEN HARDLINK TO ./usr/share/doclib/annex/man/man3/Thread.3

    Press RETURN to proceed...
```

You can disregard this information. 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.

1.1.15 Russian Keyboard (Patch 339.00)

The new Russian 3R-LKQ48–BT keyboard, for which Patch 339.00 provides an updated keyboard map, comes with five extra keycaps. To enable any of those extra keycaps, you will need to modify the file `/usr/lib/X11/xkb/symbols/digital/russian`. For example:

```
//    KEY <AD09> can be replaced by an extra keycap.
//    If you replace it with the extra keycap, please uncomment
//    the following definition and comment out the original one.
//
//    key <AD09> {
//      symbols[Group1]=3D [               o,               O ],
//      symbols[Group2]=3D [     Ukrainian_i,     Ukrainian_I ]
//    };
    key <AD09> {
        symbols[Group1]=3D [               o,               O ],
        symbols[Group2]=3D [  Cyrillic_shcha,  Cyrillic_SHCHA ]
    };
```

1.1.16 Panics May Occur on Multi-CPU Systems

Boot-time panics may occur on multi-CPU systems if all of the following conditions exist:

- Auditing is enabled.
- Audit's `-m` switch is configured to establish a dump interval.
- The system contains empty CPU slots.

The panic will occur on the first reboot after audit is configured or following an update installation on a system with audit already configured with a dump interval. The system will be unable to reboot in this configuration.

To work around this problem, boot to single-user mode and remove the `-m` option from the audit configuration stored in `/etc/rc.config.common` or `/etc/rc.config`.

The problem will be fixed in the next patch kit.

1.1.17 General and Problem Information for AlphaServer ES47, ES80, and GS1280 Systems

The following information pertains to the new AlphaServer ES47, ES80, and GS1280 systems, which require Tru64 UNIX Version 5.1B and this patch kit to be installed.

1.1.17.1 CPU Offline Restrictions

The Primary CPU cannot be taken off line.

CPUs that have I/O hoses attached to them can only be taken off line if another CPU without I/O attached is present in the system . A failure to adhere to this restriction will cause the `psradm` command to return an error.

In a two-CPU configuration, the AlphaServer ES47 and ES80 do not allow any CPUs to be taken off line.

1.1.17.2 Problem with Capacity-on-Demand Process

A problem has been discovered with the capacity on demand process in which a CPU can be designated as spare, but is not taken off line as expected.

With the capacity-on-demand process, the `codconfig` *`[cpu_id_list]`* command lets you specify which CPUs you have paid for and which are spares. The command is supposed to mark the others as spare and then take them off line. Once a CPU is marked as spare, the `hwmgr` command and Manage CPUs suitlet will not let you put them on line until you use the `ccod -l` or `ccod -p` command to either loan or purchase the CPU.

The workaround is to use the `codconfig` *`[cpu_id_list]`* command to mark the CPUs as spare, and then use either the `hwmgr` command or the Manage CPUs SUITLET to take them off line (sometimes referred to as offlining them). In the following example, *`N`* is the CPU number.

```
# hwmgr -offline -name cpuN
```

If, for example, the `codconfig` command returns the message "Error for CPU 2: Unable to offline this CPU," you would enter the following `hwmgr` command:

```
# hwmgr -offline -name cpu2
```

For more information, see `codconfig`(8) and `hwmgr`(8).

The Manage CPUs SUITLET is available from the SysMan Menu and SysMan Station.

1.1.17.3 Hardware SCSI Bus Errors

SCSI CAM errors experienced by the Adaptec controller that require SCSI bus resets could cause PCI bus faults. These faults will be seen as a “Machine Check System Uncorrectable” panic. This will require the system to be booted after the machine_check. A fix for this problem will be included in a future release.

1.1.17.4 Repeated Reboots May Cause Panic

Repeated reboots of the system may cause a kernel memory fault panic, but does not result in the loss of data. A reboot after the panic should be successful. A fix for this problem will be included in a future release.

1.1.18 Caution on Updating to Version 5.1B with DEGXA NICs

Do not attempt to do a update installation or rolling upgrade from Version 5.1A to Version 5.1B when the network device is a DEGXA-TA or DEGXA-SA and you have the Version 5.1A Patch Kit 4 and the New Hardware Devices V6 (NHD6) Kit installed.

The NHD6 kit and Patch Kit 4 have provided fixes that are not in the base operating system release for Version 5.1B. Once the update is completed using another network device and the Version 5.1B Patch Kit 1 has been applied, the DEGXA network interface cards (NICs) can again be used for the network connection.

1.1.19 Tuning the NFS Server Duplicate Request Cache (Patch 838.00)

The NFS server maintains a list of recently completed nonrepeatable requests. This list is used to reply to client retransmissions of the request in the event that the initial request transmission’s reply was lost or that the server took too long to satisfy the request.

Problems may occur with the duplicate request cache in some cases, under heavy NFS server load and over high aggregate network bandwidth involving changes to

file systems (changes caused by the use of the `creat, link, unlink, mkdir, rmdir, truncate, utimes,` and `write` commands). These problems can occur if all the elements in the duplicate request cache are cycled between an initial client transmission and subsequent retransmission. If this occurs, the NFS server cannot detect that the retransmission is in fact a retransmission. This may result in the repetition of a request and may cause out-of-order writes or truncation and subsequent retruncation of a file.

Patch 838.00 provides a tuning variable to control the size of the NFS server's duplicate request cache:

- `nfs_dupcache_size` — Controls the absolute size of the NFS server duplicate request cache. This is measured in the number of elements that are allocated at NFS server initialization.

If it is determined that the size of the duplicate cache needs to be modified, you should change `nfs_dupcache_size`. The new value for `nfs_dupcache_size` should be set to equal two times the value of `nfs_dupcache_entries`.

You must use the `dbx` command to modify `nfs_dupcache_size`. There is no sysconfig interface to this tuning variable.

1.2 Summary of Base Operating System Patches

This section provides descriptions of the patches in Patch Kit 2 for the Tru64 UNIX operating system. Because Tru64 UNIX patch kits are cumulative, each patch lists its state according to the following criteria:

- New

  Indicates a patch that is new for this release.

- New (Supersedes Patches ...)

  Indicates a patch that is new to the kit but was combined (merged) with one or more patches during the creation of earlier versions of this kit, before it was publicly released.

- Existing (Kit 1)

  Indicates a patch that was new in the previous Version 5.1B patch kit.

- Supersedes Patches ...

  Indicates a patch that was combined (merged) with other patches.

Number: Patch 2.00

Abstract: Correct improper file access

State: Existing (Kit 1)

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file access.

Number: Patch 6.00

Abstract: Correct improper file access

State: Existing (Kit 1)

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file or privilege management.

Number: Patch 8.00

Abstract: Correct improper file access

State: Existing (Kit 1)

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file or privilege management.

Number: Patch 10.00

Abstract: Correct potential buffer overflow

State: Existing (Kit 1)

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised when a buffer overflow occurs in the CDE online help. Buffer overflows are sometimes exploited in an attempt to subvert the function of a privileged program and possibly execute commands at the elevated privileges if the program file has the setuid privilege.

Number: Patch 12.00

Abstract: Correct potential buffer overflow

State: Existing (Kit 1)

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised when a buffer overflow occurs in the CDE online help. Buffer overflows are sometimes exploited in an attempt to subvert the function of a privileged program and possibly execute commands at the elevated privileges if the program file has the setuid privilege.

Number: Patch 16.00

Abstract: Fix for the dsfmgr utility

State: Existing (Kit 1)

- Fixes many small problems in dsfmgr.

Number: Patch 33.00

Abstract: Correct potential buffer overflow

State: Supersedes Patch 31.00

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised when a buffer overflow occurs in the DtSvc utility. Buffer overflows are sometimes exploited in an attempt to subvert the function of a privileged program and possibly execute commands at the elevated privileges if the program file has the setuid privilege.
- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file access.

Number: Patch 36.00

Abstract: Correct potential buffer overflow

State: Supersedes Patch 34.00

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised when a buffer overflow occurs in the DtSvc utility. Buffer overflows are sometimes exploited in an attempt to subvert the function of a privileged program and possibly execute commands at the elevated privileges if the program file has the setuid privilege.
- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file access.

Number: Patch 38.00

Abstract: Adds new functionality to tcpdump

State: Existing (Kit 1)

- Adds support for IEEE 802.1Q Virtual Local Area Network (VLAN).

Number: Patch 58.00

Abstract: Security (SSRT1-40U, SSRT1-41U, SSRT1-42U)

State: Existing (Kit 1)

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file access.

Number: Patch 60.00

Abstract: Update to siacfg utility

State: Existing (Kit 1)

- On systems using Perl 5.8.0 and higher, this patch eliminates the "Using an array as a reference is deprecated" warning when running /usr/sbin/siacfg and during system boot.

Number: Patch 62.00

Abstract: Correct improper file access

State: Existing (Kit 1)

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file or privilege management.

Number: Patch 84.00

Abstract: Security (SSRT2208)

State: Existing (Kit 1)

- Corrects a potential security vulnerability which may allow nonprivileged users to gain unauthorized (root) access. This may be in the form of local and remote security domain risks.

Number: Patch 86.00

Abstract: Security (SSRT1-40U, SSRT1-41U, SSRT1-42U)

State: Existing (Kit 1)

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file access.

Number: Patch 89.00

Abstract: Fix defects in AutoFS user space and kernel code

State: Supersedes Patch 87.00

- Fixes multiple defects in AutoFS user space and kernel code.
- Fixes a problem that prevents access to AutoFS file systems if ACLs are enabled.

Number: Patch 91.00

Abstract: Correct improper file access

State: Existing (Kit 1)

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file or privilege management.

Number: Patch 135.00

Abstract: Correct potential buffer overflow

State: Existing (Kit 1)

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised when a buffer overflow occurs in the dxterm utility. Buffer overflows are sometimes exploited in an attempt to subvert the function of a privileged program and possibly execute commands at the elevated privileges if the program file has the setuid privilege.

Number: Patch 137.00

Abstract: Fix for startslip program

State: Existing (Kit 1)

- Fixes a problem with the startslip program that prevented it from extracting all information from the acucap file.

Number: Patch 139.00

Abstract: Correct improper file access

State: Existing (Kit 1)

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file access.

Number: Patch 157.00

Abstract: Correct improper file access

State: Existing (Kit 1)

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file or privilege management.

Number: Patch 159.00

Abstract: Fix hwmgr command to show path state correctly

State: Existing (Kit 1)

- Fixes the hwmgr command to correctly show a path state.

Number: Patch 167.00

Abstract: Security (SSRT2368, SSRT2368)

State: Existing (Kit 1)

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper access.

Number: Patch 169.00

Abstract: Fix for SDLT media error

State: Existing (Kit 1)

- Adds the capability for KZPCA devices to work with SCSI devices that only support asynchronous data transfers and fixes SDLT media error caused bus resets.

Number: Patch 171.00

Abstract: Correct improper file access

State: Existing (Kit 1)

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file access.

Number: Patch 173.00

Abstract: Fix race condition and improper file access

State: Existing (Kit 1)

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file access.

Number: Patch 176.00

Abstract: Resolve kernel memory faults in TCP/IP subsystem

State: Supersedes Patch 174.00

- Resolves kernel memory faults in the TCP/IP subsystem.

Number: Patch 178.00

Abstract: Fix for lpd line printer daemon

State: Existing (Kit 1)

- Fixes the lpd daemon to correct /etc/hosts.lpd case sensitivity, for example, “node.domain" treated the same as "Node.Domain/"

Number: Patch 185.00

Abstract: Prevents addvol from adding invalid disks into domain

State: Existing (Kit 1)

- Prevents addvol from adding invalid disks into a domain.

Number: Patch 187.00

Abstract: Fix for invalid disks being added into domain

State: Existing (Kit 1)

- Prevents addvol from adding invalid disks into a domain.

Number: Patch 191.00

Abstract: Correct improper file access

State: Existing (Kit 1)

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file access.

Number: Patch 193.00

Abstract: Correct improper file access

State: Existing (Kit 1)

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file access.

Number: Patch 197.00

Abstract: Read privileges being stripped from passwd file

State: Existing (Kit 1)

- Fixes a problem in which group and other read privileges get stripped from /etc/passwd when a user switches from enhanced to base security.

Number: Patch 199.00

Abstract: Correct improper file access

State: Existing (Kit 1)

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file or privilege management.

Number: Patch 201.00

Abstract: Fix for verify utility

State: Existing (Kit 1)

- Fixes a problem in which the verify utility core dumps if it encounters a specific type of metadata inconsistency.

Number: Patch 203.00

Abstract: Correct improper file access

State: Existing (Kit 1)

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file access.

Number: Patch 205.00

Abstract: Fixes a correctable error reporting problem

State: Existing (Kit 1)

- Fixes a correctable error reporting problem that turns off the reporting of correctable errors forever on any CPU, except CPU 0, once throttling of correctable errors has begun.

Number: Patch 207.00

Abstract: Correct potential buffer overflow

State: Existing (Kit 1)

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised when a buffer overflow occurs in the libXm utility. Buffer overflows are sometimes exploited in an attempt to subvert the function of a privileged program and possibly execute commands at the elevated privileges if the program file has the setuid privilege.

Number: Patch 209.00

Abstract: Correct potential buffer overflow

State: Existing (Kit 1)

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised when a buffer overflow occurs in the libXm utility. Buffer overflows are sometimes exploited in an attempt to subvert the function of a privileged program and possibly execute commands at the elevated privileges if the program file has the setuid privilege.

Number: Patch 219.00

Abstract: Correct improper file access

State: Existing (Kit 1)

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file or privilege management.

Number: Patch 221.00

Abstract: Correct improper file access

State: Existing (Kit 1)

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file access.

Number: Patch 223.00

Abstract: Add SCSI reserve/release support to mt

State: Existing (Kit 1)

- Adds SCSI reserve and release support to the mt program to assist open SAN tape management.

Number: Patch 225.00

Abstract: Interop problem between curses.h and esnmp.h.

State: Existing (Kit 1)

- Fixes an interoperability problem between the curses.h and esnmp.h header files.

Number: Patch 229.00

Abstract: Correct improper file access

State: Existing (Kit 1)

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file or privilege management.

Number: Patch 232.00

Abstract: Correct improper file access

State: Supersedes Patch 230.00

- Adds support in script to remove all Persistant Reservations for MSA controller.
- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file access.

Number: Patch 234.00

Abstract: Correct improper file access

State: Existing (Kit 1)

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file or privilege management.

Number: Patch 238.00

Abstract: Correct improper file access

State: Existing (Kit 1)

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file or privilege management.

Number: Patch 240.00

Abstract: Fix for hwmgr command

State: Existing (Kit 1)

- Fixes a problem in which the display for the hwmgr -show name command is not aligned properly for the name field.

Number: Patch 248.00

Abstract: Fix for hwmgr delete command option

State: Existing (Kit 1)

- Fixes a problem where, when using hwmgr to delete a component, a "DELETE_COMMIT: Cannot fetch name." message may be displayed on the console. This problem can be seen frequently in a cluster environment when the component being deleted does not exist on the system.

Number: Patch 250.00

Abstract: Correct improper file access

State: Existing (Kit 1)

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file access.

Number: Patch 254.00

Abstract: Correct improper file access

State: Existing (Kit 1)

- Corrects the problem of a core dump that occurs when the output from the lint program for a nonexistent file is supplied to error.
- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file or privilege management.

Number: Patch 262.00

Abstract: Correct improper file access

State: Existing (Kit 1)

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file access.

Number: Patch 264.00

Abstract: Correct improper file access

State: Existing (Kit 1)

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file access.

Number: Patch 266.00

Abstract: Correct improper file access

State: Existing (Kit 1)

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file access.

Number: Patch 268.00

Abstract: Installs version V2.1-120 of libots3 libraries

State: Existing (Kit 1)

- Installs version V2.1-120 of /usr/lib/libots3.a and /usr/shlib/libots3.so, which fixes a problem where long-running OpenMP applications might overflow an internal libots3 counter, resulting in a breakdown of thread synchronization.

Number: Patch 270.00

Abstract: Installs version V2.1-120 of libots3 libraries

State: Existing (Kit 1)

- Installs version V2.1-120 of /usr/lib/libots3.a and /usr/shlib/libots3.so, which fixes a problem where long-running OpenMP applications might overflow an internal libots3 counter, resulting in a breakdown of thread synchronization.

Number: Patch 272.00

Abstract: Correct improper file access

State: Existing (Kit 1)

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file access.

Number: Patch 275.00

Abstract: Fix for hwmgr -view transaction -cluster command

State: Supersedes Patch 273.00

- Corrects some command-parsing irregularities in hwmgr that may cause options like -category and -cluster to be confused.
- Corrects a problem in which information from the hwmgr -view transaction -cluster command for a node on a cluster may not be displayed.

Number: Patch 279.00

Abstract: Corrections to Oxygen VX1 graphics card XCopyPlane

State: Existing (Kit 1)

- Corrects a problem with the Oxygen VX1 graphics card to make XCopyPlane copy only the requested bitplane rather than all bitplanes.

Number: Patch 281.00

Abstract: Fixes a problem in usb_hid.mod

State: Existing (Kit 1)

- Corrects a problem in which a kernel memory fault sometimes occurs if a USB keyboard or mouse does not respond quickly enough. This KMF can occur during boot or soon after a USB keyboard or mouse is connected. Any device can trigger this, though it is neither predictable nor common.

Number: Patch 283.00

Abstract: Correct improper file access

State: Existing (Kit 1)

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file access.

Number: Patch 285.00

Abstract: Correct improper file access

State: Existing (Kit 1)

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file access.

Number: Patch 289.00

Abstract: Correct improper file access

State: Existing (Kit 1)

- Corrects a potential security vulnerability which, under certain circumstances, could compromise system integrity.

Number: Patch 296.00

Abstract: Correct improper file access

State: Existing (Kit 1)

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file access.

Number: Patch 298.00

Abstract: Fixes consvar -s bootdef_dev failure with KZPCC

State: Existing (Kit 1)

- Fixes consvar -s bootdef_dev failure with KZPCC.

Number: Patch 300.00

Abstract: Login process crashes when LDAP users try to log in

State: Existing (Kit 1)

- Fixes a problem in with the login process may crash when LDAP users or users belonging to an LDAP group attempt to log in.

Number: Patch 306.00

Abstract: Correct improper file access

State: Existing (Kit 1)

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file access.

Number: Patch 309.00

Abstract: Performance tool failures on Sierra Cluster (PFS)

State: Supersedes Patch 307.00

- Fixes a problem in which the prof -pixie -testcoverage <exe> <exe>.Counts sometimes reports invalid source line number ranges.
- Fixes performance tool failures on Sierra Clusters Parallel File Systems (PFS) .

Number: Patch 311.00

Abstract: Correct improper file access

State: Existing (Kit 1)

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file access.

Number: Patch 313.00

Abstract: Correct improper file access

State: Existing (Kit 1)

- Corrects a potential security vulnerability where under certain circumstances, system integrity may be compromised. This may be in the form of improper file or privilege management.
- Addresses a problem in which performing a sort on a large database using numerous keys fails during the consolidation phase of the temporary files.

Number: Patch 317.00

Abstract: Security (SSRT1-40U, SSRT1-41U, SSRT1-42U)

State: Existing (Kit 1)

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file access.

Number: Patch 320.00

Abstract: Fixes problem with IPv6 neighbor discovery daemon

State: Supersedes Patch 318.00

- Fixes a regression in the operation of the IPv6 neighbor discovery daemon, where IPv6 addresses will not be automatically configured on PPP interfaces.
- Fixes a problem with IPv6 neighbor discovery daemon, where under certain circumstances, the daemon can cause bad information to be written to a DNS database, thereby causing failures on subsequent database reloads.

Number: Patch 322.00

Abstract: Correct improper file access

State: Existing (Kit 1)

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file or privilege management.

Number: Patch 324.00

Abstract: Correct improper file access

State: Existing (Kit 1)

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file or privilege management.

Number: Patch 326.00

Abstract: Fixes a linker error

State: Existing (Kit 1)

- Fixes a linker error that occurs when running the command ld -update_registry /dev/null.

Number: Patch 328.00

Abstract: Fixes and improves the mcutil program

State: Existing (Kit 1)

- Fixes and improves the mcutil program by correcting how bus resets are handled by the program and enhancing its error reporting capabilities.

Number: Patch 330.00

Abstract: Allows evmd to stop listening on default TCP port 619

State: Existing (Kit 1)

- Allows the Event Manager daemon, evmd, to stop listening on its default TCP port 619. This capability is not available for clustered systems.

Number: Patch 332.00

Abstract: Fixes memory leak in the PanoramiX/Xinerama Extension

State: Existing (Kit 1)

- Fixes a memory leak in the PanoramiX/Xinerama Extension that could cause a process core dump.

Number: Patch 334.00

Abstract: Correct improper file access

State: Existing (Kit 1)

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file access.

Number: Patch 339.00

Abstract: Updated keyboard map for Russian 3R-LKQ48-BT

State: Existing (Kit 1)

- Provides an updated keyboard map for the Russian 3R-LKQ48-BT keyboard model.

Number: Patch 343.00

Abstract: Fixes problem seen with TAHI IPv6 conformance test

State: Existing (Kit 1)

- Fixes a problem seen with the TAHI IPv6 conformance test, specifically Test 4 for the IPv6 Specification.

Number: Patch 345.00

Abstract: Correct improper file access

State: Existing (Kit 1)

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file or privilege management.

Number: Patch 347.00

Abstract: Fix allows fuser to display the reference flag

State: Existing (Kit 1)

- Allows the fuser utility to display the reference flag, which indicates the type of reference made; for example, open, closed, unlinked, or mmapped.

Number: Patch 351.00

Abstract: Correct improper file access

State: Existing (Kit 1)

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file access.

Number: Patch 356.00

Abstract: Fix for ftp open command

State: Supersedes Patch 354.00

- Corrects a bug in the ftp open command. The optional port argument now accepts port numbers between 32768 and 65535.
- Corrects a potential security vulnerability which, under certain circumstances, could compromise system integrity.

Number: Patch 358.00

Abstract: Fix for mountd daemon

State: Existing (Kit 1)

- Enables mountd to correctly handle entries with multiple lines input in exports file.

Number: Patch 360.00

Abstract: Correct improper file access

State: Existing (Kit 1)

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file or privilege management.

Number: Patch 362.00

Abstract: Correct improper file access

State: Existing (Kit 1)

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file access.

Number: Patch 364.00

Abstract: Correct improper file access

State: Existing (Kit 1)

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file or privilege management.

Number: Patch 366.00

Abstract: Correct improper file access

State: Existing (Kit 1)

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file access.

Number: Patch 368.00

Abstract: Correct improper file access

State: Existing (Kit 1)

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file access.

Number: Patch 370.00

Abstract: Correct improper file access

State: Existing (Kit 1)

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file or privilege management.

Number: Patch 372.00

Abstract: Security (SSRT1-40U, SSRT1-41U, SSRT1-42U)

State: Existing (Kit 1)

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file access.

Number: Patch 378.00

Abstract: Security (SSRT1-40U, SSRT1-41U, SSRT1-42U)

State: Existing (Kit 1)

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file access.

Number: Patch 380.00

Abstract: Correct improper file access

State: Existing (Kit 1)

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file or privilege management.

Number: Patch 382.00

Abstract: Fixes a simple lock panic in the floppy driver

State: Existing (Kit 1)

- Fixes a simple lock panic in the floppy driver.

Number: Patch 384.00

Abstract: Fix for telnetd daemon

State: Existing (Kit 1)

Fixes a problem in which telnetting from some machines (MS), will leave a UDP port open, requiring a call to yp_unbind() after getnameinfo() to close all ports.

Number: Patch 387.00

Abstract: Correct improper file access

State: Existing (Kit 1)

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file or privilege management.

Number: Patch 391.00

Abstract: Updates reference pages for VLAN functionality

State: Existing (Kit 1)

- Revises the ifconfig(8), lan_config(8), niffconfig(8), ping(8), vlanconfig(8), and vlan(7) reference pages for VLAN functionality.

Number: Patch 396.00

Abstract: Fix for ldapcd daemon

State: Supersedes Patch 56.00

- Fixes the following problems with the ldapcd daemon:
  - It may crash when resolving group codes with very large GIDs.
  - It may crash when the LDAP Directory Server is unavailable.
  - It prevents LDAP users from being authenticated, even when they are providing the correct password.

Number: Patch 414.00

Abstract: Revises several SSH reference pages

State: Existing (Kit 1)

- Revises several of the SSH reference pages which address several issues and problems with SSH, including the following:
  - Interoperability with other SSH implementations
  - Client/server configuration files compatibility issues
  - The lack of IPV6 support

Number: Patch 416.00

Abstract: Fix for creacct hang

State: Existing (Kit 1)

- Fixes a problem that causes the creacct command to hang when the W2K Active Directory is misconfigured.

Number: Patch 421.00

Abstract: Revises envconfig.8 and envmond.8 reference pages

State: Existing (Kit 1)

- Revises the envconfig(8) and envmond(8) reference pages for the environmental monitoring facilities /usr/sbin/envmond and /usr/sbin/envconfig to support the new GS1280 hardware platform.

Number: Patch 423.00

Abstract: Fix potential denial of service

State: Existing (Kit 1)

- Corrects a potential security vulnerability for systems using Internet Protocol Security (IPsec). Under certain circumstances, a remote attacker may be able to cause IPsec to block all IP traffic from the system, creating a denial of service.

Number: Patch 433.00

Abstract: Fix race condition and improper file access

State: Supersedes Patch 236.00

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file access.
- Corrects a problem in which a core dump may occur when using the csh shell from the Japanese locale.
- Fixes the problem with csh shell redirection that occurs while using a tilde (~) operation to redirect standard input and standard output of a command to a file residing in the home directory.

Number: Patch 442.00

Abstract: Incorrect I/O status may be returned by KZPEA driver

State: Supersedes Patch 315.00

- Corrects problems in the aha_chim driver that could result in bus hangs, panics, and inappropriate access of freed memory during high rate of bus resets.
- Corrects a problem in which Incorrect I/O status may be returned by the KZPEA driver when attempting to abort an I/O during a reset.

Number: Patch 444.00

Abstract: Revises tcpdump.8 ref page for VLAN functionality

State: Existing (Kit 1)

- Revises the tcpdump(8) reference page for virtual local area network (VLAN) functionality.

Number: Patch 446.00

Abstract: Revises the mt.1 reference page

State: Existing (Kit 1)

- Revises the mt(1) reference page for the mt command, which has three new commands, mt reserve, mt release and mt tur.

Number: Patch 448.00

Abstract: Allows multiple VX1 graphic cards to be configured

State: Supersedes Patch 353.00

- Corrects a problem in which systems configured with VX1 graphics card will not return to console when the halt button is pressed, thereby making the console unusable.
- Allows multiple VX1 graphic cards to be configured in a separate I/O box system.

Number: Patch 452.00

Abstract: Modifications for environmental monitoring facilities

State: Supersedes Patch 40.00

- Modifies the environmental monitoring facilities /usr/sbin/envmond and /usr/sbin/envconfig to support the AlphaServer GS1280 system.
- Updates the environmental monitoring daemon envmond to ensure the correct EVM events are being sent at the correct time.

Number: Patch 454.00

Abstract: Correct potential buffer overflow

State: Existing (Kit 1)

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised when a buffer overflow occurs in the dxsysinfo utility. Buffer overflows are sometimes exploited in an attempt to subvert the function of a privileged program and possibly execute commands at the elevated privileges if the program file has the setuid privilege.

Number: Patch 458.00

Abstract: Correct potential buffer overflow

State: Existing (Kit 1)

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file or privilege management.

Number: Patch 476.00

Abstract: Corrects hang in log command

State: Existing (Kit 1)

- Corrects possible deadlock in the ./isl/log and ./usr/sbin/log commands.

Number: Patch 486.00

Abstract: Security (SSRT0785U)

State: Supersedes Patches 162.00, 163.00, 165.00

- Corrects several problems with the account management tools, including the following:
  - The userdel command possibly core dumping when the shell field is empty in the passwd file.
  - The usermod command not working as expected with NIS +/- users.
  - The useradd command not managing default and template data properly. This showed up most notably with the useradd -p command producing the message "Password must be between 32 and 80 characters."
- Updates the account management tools to use the latest versions of the ASU (Advanced Server for UNIX) API calls when ASU is in use on the server.
- Fixes a number of problems with the Account Manager application, dxaccounts, on a system with ASU installed.
- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of passwords that have a length outside of the intended range.

Number: Patch 492.00

Abstract: Correct improper file access

State: Supersedes Patch 77.00

- Corrects several potential security vulnerabilities where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file access.

Number: Patch 502.00

Abstract: Fixes Solaris rcp commands failures

State: New

- Corrects a problem in which Solaris rcp commands would fail against Tru64 UNIX V5.1B servers with the message "rcp: lost connection".

Number: Patch 506.00

Abstract: Revised newfs(8) reference page

State: New

- Revises the newfs(8) reference page to document a -M command option to newfs to allow users to specify permissions of an mfs root directory when it is first created.

Number: Patch 525.00

Abstract: Security (SSRT2301, SSRT2275)

State: Supersedes Patch 20.00

- Provides protection against several potential security vulnerabilities where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file access.
- Provides protection against a class of potential security vulnerabilities called buffer overflows. Buffer overflows are sometimes exploited in an attempt to subvert the function of a privileged program and possibly execute commands at the elevated privileges if the program file has the setuid privilege. This patch allows a system administrator to enable memory management protections that limit potential buffer overflow vulnerabilities.
- Provides protection against a potential security vulnerability, where under certain circumstances, system integrity may be compromised when a buffer overflow occurs in the uucp utility.

Number: Patch 527.00

Abstract: Security (SSRT2275)

State: Supersedes Patch 258.00

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file or privilege management.
- Provides protection against a class of potential security vulnerabilities called buffer overflows. Buffer overflows are sometimes exploited in an attempt to subvert the function of a privileged program and possibly execute commands at the elevated privileges if the program file has the setuid privilege. This patch allows a system administrator to enable memory management protections that limit potential buffer overflow vulnerabilities.

Number: Patch 529.00

Abstract: Security (SSR)T2275

State: New

- Provides protection against a class of potential security vulnerabilities called buffer overflows. Buffer overflows are sometimes exploited in an attempt to subvert the function of a privileged program and possibly execute commands at the elevated privileges if the program file has the setuid privilege. This patch allows a system administrator to enable memory management protections that limit potential buffer overflow vulnerabilities.

Number: Patch 531.00

Abstract: Security (SSRT2275)

State: New

- Provides protection against a class of potential security vulnerabilities called buffer overflows. Buffer overflows are sometimes exploited in an attempt to subvert the function of a privileged program and possibly execute commands at the elevated privileges if the program file has the setuid privilege. This patch allows a system administrator to enable memory management protections that limit potential buffer overflow vulnerabilities.

Number: Patch 533.00

Abstract: Security (SSRT2275)

State: New

- Provides protection against a class of potential security vulnerabilities called buffer overflows. Buffer overflows are sometimes exploited in an attempt to subvert the function of a privileged program and possibly execute commands at the elevated privileges if the program file has the setuid privilege. This patch allows a system administrator to enable memory management protections that limit potential buffer overflow vulnerabilities.

Number: Patch 535.00

Abstract: Corrects problems in dbx and object file tools

State: New

- Fixes various problems in the dbx and object file tools: dbx, ostrip, strip, mcs, dis, cord, file, and stdump.

Number: Patch 537.00

Abstract: Fixes various problems in libc functions

State: New

- Fixes various problems in the libc functions getdate(), strptime(), callrpc(), strncasecmp() and fork(). It also fixes a problem in the libnuma function nacreate() and the system header <sgtty.h>.

Number: Patch 539.00

Abstract: Fixes client hangs under Enhanced Security

State: New

- Fixes client (login, su, rshd, edauth, and sshd2) hangs and long delays under Enhanced Security, as well as some intermittent errors or failures seen with prpasswdd or rpc.yppasswdd.

Number: Patch 541.00

Abstract: Fixes various problems in libc functions

State: New

- Fixes various problems in the libc functions getdate(), strptime(), callrpc(), strncasecmp() and fork().
- Fixes a problem in the libnuma function nacreate() and the system header <sgtty.h>.

Number: Patch 543.00

Abstract: Update to prpasswdd daemon

State: New

- Fixes client (login, su, rshd, edauth, and sshd2) hangs and long delays under Enhanced Security, as well as some intermittent errors or failures seen with prpasswdd or rpc.yppasswdd.

Number: Patch 545.00

Abstract: Fixes problems in dbx and object file tools

State: New

- Fixes various problems in the dbx and object file tools: dbx, ostrip, strip, mcs, dis, cord, file, and stdump.

Number: Patch 547.00

Abstract: Fixes problem in XTI caused by blocked mutex lock

State: New

- Fixes a problem in XTI caused by a blocked mutex lock that caused any thread attempting to send an abortive disconnect to hang.

Number: Patch 549.00

Abstract: Fixes problem in XTI caused by blocked mutex lock

State: New

- Fixes a problem in XTI caused by a blocked mutex lock that caused any thread attempting to send an abortive disconnect to hang.

Number: Patch 551.00

Abstract: Corrects improper file or privilege

State: New

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file or privilege management.

Number: Patch 553.00

Abstract: Corrects improper file or privilege management

State: Supersedes Patches 304.00, 189.00

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file or privilege management.
- Corrects a problem found wherein the rmtmpfiles script would leave empty directories in /var/tmp at system startup.
- Makes startup scripts in /sbin/init.d world readable.

Number: Patch 555.00
Abstract: Scripts in /sbin/init.d are now world readable
State: New

- Makes startup scripts in /sbin/init.d world readable.

Number: Patch 557.00
Abstract: Scripts in /sbin/init.d are now world readable
State: New

- Makes startup scripts in /sbin/init.d world readable.

Number: Patch 560.00
Abstract: Update to nis startup script
State: New (Supersedes Patch 558.00)

- Fixes a typo in mkcdsl.
- Updates the NIS startup script to correctly start NIS on the cluster alias.
- Makes startup scripts in /sbin/init.d world readable.

Number: Patch 562.00
Abstract: Scripts in /sbin/init.d are now world readable
State: New

- Makes start-up scripts in /sbin/init.d world readable.

Number: Patch 564.00
Abstract: Correct improper file access
State: Supersedes Patches 335.00, 337.00, 392.00, 394.00

- Corrects a problem in which volmigrate returns a shell error when attempting to migrate an AdvFS domain with multiple filesets. With this patch, the domains can be migrated if all the filesets are mounted.
- Corrects several potential security vulnerabilities where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file access.
- Prevents inconsistent LSM volumes when the name of a partition that is being encapsulated matches the name of a current LSM volume.

Number: Patch 566.00
Abstract: Fixes problems in the kdbx extensions
State: Supersedes Patch 4.00

- Fixes a premature termination of the ofile kdbx extension and warning messages in various kdbx extensions.
- Fixes problems in the kdbx u and vnode extensions

Number: Patch 576.00
Abstract: Correct improper file access
State: Supersedes Patch 47.00

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file or privilege management.
- Removes files created under /usr/lib/sabt/ during the running of the btcreate utility on file systems with LSM. Previously, these files were copied to the tape, and would be restored as if they were archived by btcreate.

Number: Patch 578.00
Abstract: Fix for grep command hang
State: New

- Changes the grep command to allow blank lines in the pattern file, and prevents it from hanging when executed with -w and -f options.

Number: Patch 622.00

Abstract: Hang occurs with various PowerStorm graphics options

State: Supersedes Patch 107.00

- Fixes a problem where the X server's command line option to turn off VESA Display Power Management Signalling (-dpms) does not work.
- Corrects a problem in which the X server may hang every 49 days on systems with PowerStorm 4D40T, 4D50T, 4D51T, or 4D60T graphics options.

Number: Patch 637.00

Abstract: Fix for .mrg..termcap file

State: New

- Corrects a problem with the merging of the termcap file during a rolling upgrade.

Number: Patch 640.00

Abstract: Fixes hang in the hardware configuration subsystem

State: Supersedes Patches 144.00, 145.00, 146.00, 147.00, 148.00, 149.00, 151.00, 484.00, 638.00

- Corrects a problem where after entering a hwmgr -redirect scsi command and rebooting, the system only boots to single user mode with the following error displayed: bcheckrc:

  Device Naming failed boot configure or verify Please correct the problem and continue or reboot INIT: SINGLE-USER MODE
- Corrects a problem in which incorrect values for LONG_MAX and LONG_MIN were displayed when using the hardware manager to show attributes.
- Addresses a problem encountered when mounting cluster root if the cluster root domain devices are private to different cluster members. Currently, you cannot boot your cluster. It will hang. With this fix, your cluster will boot with a warning to the console. This configuration is not recommended; however the cluster should not be unbootable. Currently, this is with respect to non-LSM cluster root domains.
- Introduces type checking of attributes when registering components with the hardware manager.
- Corrects a potential deadlock in the hardware configuration subsystem.
- Prevents the hardware management cluster database from being reset.
- Corrects invalid hwmgr show component inconsistency.
- Corrects a problem encountered when hwmgr is used to verify that a sensor's status would change from OK to Fault; each time the state changed and hwmgr requested the new value, hwmgr dumped core.
- Suppresses an erroneous console warning message that may be provided when cluster root is under LSM control. The warning "WARNING: cluster root devices are on private buses!" may be erroneously output when cluster root is under LSM control. LSM does not support such configurations.
- Fixes a possible, but rare, hang in the hardware configuration subsystem.

Number: Patch 643.00

Abstract: Fix for tar command

State: Supersedes Patches 105.00, 641.00

- Fixes a one-byte gap/hole in the maximum file size in the tar command before an extended header record is used (8589934591 (octal 77777777777)).
- Makes the following changes to the tar, pax, and cpio commands:
  - Tar now checks and reports any write errors.
  - The tar, pax, and cpio commands can unalter the ctime of input files upon creation of an archive and display a warning message if unable to preserve the time of input files.
  - Corrects the tar o option behavior.
  - Corrects the pax -l option to create hard links properly.
  - Corrects the cpio -o option to not corrupt extended UID file ownership.
  - Fixes how long file names are handled in tar.
  - Fixes pax to handle ACL on directories properly.
  - Corrects tar to properly handle unusual directory specifications.

Number: Patch 646.00

Abstract: Correct improper file access

State: Supersedes Patches 227.00, 644.00

- Fixes a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file or privilege management.
- Corrects the following problems found in accounting commands:
  - Fixes the way accounting files are referenced using CDSLs.
  - Resolves the differences in the CPU time and connect time found during the conversion of accounting reports from ASCII format to binary and again back to ASCII.
  - Resolves the differences in CPU time found in the output of the acctcom and acctmerg commands for the same input file.

Number: Patch 648.00

Abstract: Fix for fwtmp command

State: New

- Fixes the fwtmp command so it does not display the invalid (negative) PIDs when the number of decimal digits of PID value exceeds 5.

Number: Patch 650.00

Abstract: Correct potential buffer overflow

State: New

Addresses potential BIND (Berkeley Internet Name Domain) security vulnerabilities that may result in buffer overflows, unauthorized access, or denial of service. These potential security vulnerabilities may be in the form of local and remote security domain risks. The following potential security vulnerabilities have been corrected:

SSRT2408 BIND - (Severity - High)
SSRT2410 BIND - (Severity - High)
SSRT2411 BIND - (Severity - High)

Number: Patch 653.00

Abstract: Correct improper file or privilege

State: New (Supersedes Patch 651.00)

- Addresses a problem of compiler warnings caused by calling a function with too few arguments.
- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file or privilege management.

Number: Patch 656.00

Abstract: Supports monitoring AdvFS statistics using pmgrd

State: New (Supersedes Patches 654.00)

- Provides support for monitoring AdvFS statistics using the Performance Manager metrics server daemon, pmgrd.
- Provides support for monitoring the disk I/O rate using pmgrd.

Number: Patch 658.00

Abstract: Fix for rm command

State: New

- Addresses a performance issue of rm -r with large directories.

Number: Patch 662.00

Abstract: mtools now prints appropriate error messages

State: New

- Modifies the mformat prompt.
- Changes mcopy and mwrite so they can overwrite existing files.
- Changes mtools so they can print appropriate error messages.
- Addresses security vulnerabilities.

Number: Patch 664.00

Abstract: Fixes erroneous reboot of the operating system

State: New

- Fixes a situation where, on a panic, the OS will erroneously reboot instead of halt and fail to take a crash dump.

Number: Patch 667.00

Abstract: Security (SSRT0711U)

State: Supersedes Patches 211.00, 429.00, 665.00

- Fixes a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file or privilege management.
- Makes the following changes to cron:
  - Corrects a problem in which crontab removes its entries and the vi editor truncates an existing file when a file system is full.
  - Corrects the improper scheduling of cron jobs related to months not having 31 days.
  - Enhances cron to now do extensive logging.

Number: Patch 669.00

Abstract: Added sysconfig tunables for ES45 environ monitoring

State: Supersedes Patches 214.00, 215.00, 217.00, 431.00

- Fixes a problem that can cause an AlphaServer ES45 system to hang if the Xserver is restarted or the system rebooted without a power cycle when using the Radeon AGP graphics device.
- Prevents the memory troller from starting on titan and tsunami platforms with aluminum ev68 CPUs.
- Fixes several IPMI-related problems, including the following:
  - Erroneous fields in 686 OS-detected environmental machine check logout frame
  - Unusually large number of 686 sensor timeouts with heavy system load
  - IPMI always reporting -48v sensors as broken, seen as "redundant power supply failed" messages
  - An IPMI memory leak
- Provides additional environmental support functionality for the AlphaServer DS20L system.
- Adds sysconfig tunables for ES45 environmental monitoring.

Number: Patch 671.00

Abstract: Fix for mfs command

State: New

- Adds a -M option to the newfs command that allows users to specify permissions of an mfs root directory when it is first created.
- Adds EVM notification support for UFS file systems.

Number: Patch 674.00

Abstract: Fix for ln command

State: New (Supersedes Patch 672.00)

- Eliminates compiler warnings in ln.
- Corrects the behavior of ln -sf to address the issue caused when a symbolic link points to a nonexisting file.

Number: Patch 678.00

Abstract: Fix race condition and improper file access

State: Supersedes Patches 143.00, 675.00, 676.00

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file access.
- Eliminates compiler warnings in ksh.
- Fixes a ksh problem related to cleaning the process when a terminal is abruptly stopped.

Number: Patch 680.00

Abstract: Fix for audit_tool command

State: Supersedes Patch 252.00

- Changes the audit_tool command as follows:
  - Fixes the search algorithm to differentiate between prived and non-prived UIDS, and to allow regular expressions in string searches.
  - Fixes a problem in which nonsense characters are appended to the audit information output of an execve event in brief mode.

Number: Patch 685.00

Abstract: Fix for vrestore command

State: New (Supersedes Patches 681.00, 682.00, 683.00)

- Makes the following changes to the vdump and vrestore commands:
  - Causes vdump and vrestore to act as expected upon receiving an interrupt (^C).
  - Fixes vdump and vrestore to pick up correct messages in all locales.
  - Increases the maximum blocksize for vdump and vrestore (performance).
  - Causes vdump to avoid some unnecessary function calls, thereby allowing faster vdumps.
  - Fixes vrestore to display bit file attributes with the -l option.
  - Prevents vrestore from failing during a remote system call.
  - Causes vrestore to display a file and directory name along with the error message when the command fails to set a property list.
  - Prevents vrestore from dumping core when when a tape has a smaller blocksize than expected.
  - Allows vrestore to handle no-rewind tapes properly.
  - Lets vrestore read environment variables for a user-defined device name.
  - Allows attributes to set to the top level-directory.

Number: Patch 687.00

Abstract: Removes compiler warnings

State: New

- Removes compiler warnings addressing outside of array bounds.

Number: Patch 689.00

Abstract: Corrects a problem in os_mibs

State: New

- Corrects a problem in os_mibs that results in the swap size and swap used values for the host mib being reported as negative values on some systems.

Number: Patch 692.00

Abstract: Fix for dump command

State: New (Supersedes Patch 690.00)

- Fixes the dump command to recognize LSM volumes correctly and to not report random information when an error has occurred.
- Introduces dumprmt.msg for remote dump/restore messages. This new message catalog file is used in both rdump and rrestore programs.

Number: Patch 694.00

Abstract: Corrects exit status of sed when disk is full

State: New

- Corrects the exit status of sed when the disk is full.

Number: Patch 696.00

Abstract: Revises the sys_attrs_ee(5) reference page

State: New

- Revises the sys_attrs_ee(5) reference page to document the new ee subsystem attribute link_check_interval.

Number: Patch 698.00

Abstract: Correct improper file access

State: Supersedes Patch 287.00

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file access.
- Addresses problems with the mksas utility where false warning messages are generated and where the user-specified temporary directory could be erroneously removed.

Number: Patch 703.00

Abstract: Correct improper file or privilege management

State: New

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file or privilege management.

Number: Patch 705.00

Abstract: Eliminates compiler warnings in mkdir

State: New

- Eliminates compiler warnings in the mkdir command.

Number: Patch 708.00

Abstract: Correct potential buffer overflow

State: Supersedes Patches 460.00, 706.00

- Corrects several potential security vulnerabilities where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file or privilege management.
- Fixes the message catalog for the CDE application dtprintinfo.

Number: Patch 710.00

Abstract: Fix for ps command

State: New

- Allows whitespace in header field and in multiple headers with the ps option -o.

Number: Patch 712.00

Abstract: Provides the chatr(1) reference page

State: New

- Adds the chatr(1) reference page for the chatr command.

Number: Patch 714.00

Abstract: Fix for vmstat

State: New

- Corrects a problem in which vmstat may display incorrect free page counts on NUMA systems.

Number: Patch 716.00

Abstract: Performance enhancement for IPsec

State: New

Corrects a condition in which IPsec performs poorly when using ESP with the 3DES cipher.

Number: Patch 718.00

Abstract: Provides correct labels for mach events

State: New

- Provides the correct labels to the audit subsystem for mach events.

Number: Patch 720.00

Abstract: Fix for audit subsystem utilities

State: New

- Provides the correct labels to the audit subsystem for mach events.

Number: Patch 722.00

Abstract: Fixes a typo in mkcdsl

State: New

- Fixes a typo in mkcdsl.
- Updates the NIS startup script to correctly start NIS on the cluster alias.

Number: Patch 727.00

Abstract: Fix for which command

State: New

- Fixes /usr/bin/which to take path information from environment rather ~/.cshrc if it is invoked from other than the C shell.

Number: Patch 729.00

Abstract: Correct improper file or privilege management

State: New

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file or privilege management.

Number: Patch 731.00

Abstract: Update to pmgrd IoRate Statistics feature

State: New

- Adds a new table in pm.mib for pmgrd IoRate Statistics feature.

Number: Patch 733.00

Abstract: Correct improper file or privilege management

State: New

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file or privilege management.

Number: Patch 735.00

Abstract: Fix for chfile command

State: New

- Causes the chfile command to display an nformative error message if it fails while trying to enable data logging.

Number: Patch 737.00

Abstract: Fix for tip command

State: New

- Enables the tip command to log into the member-specific log file.
- Corrects the path for the aculog file and gives appropriate permissions.

Number: Patch 739.00

Abstract: Enables tip to log into member-specific log file

State: New

- Enables the tip command to log into the member-specific log file.
- Corrects the path for the aculog file and gives appropriate permissions.

Number: Patch 741.00

Abstract: Fix for default cron jobs

State: New

- Causes the rescheduling of certain default cleanup cron jobs so that they will not get skipped during a time change to daylight savings time (DST).

Number: Patch 743.00

Abstract: Corrects a problem in niffd

State: New

- Corrects a problem with the Network Interface Failure Finder daemon, niffd, that causes its memory usage to grow over time.

Number: Patch 745.00

Abstract: Correct improper file or privilege management

State: New

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file or privilege management.

Number: Patch 747.00

Abstract: Fix for btextract utility

State: New

- Corrects a problem in which the btextract utility was not preventing the advanced mode of restore for a system with LSM setup.

Number: Patch 749.00

Abstract: make command now checks dependencies on archive libraries

State: New

- Changes the /usr/opt/ultrix/usr/bin/make command to properly check dependencies on archive libraries.

Number: Patch 751.00

Abstract: Fix for .mrg...login script

State: New

- Adds informative messages during a rolling upgrade when a problem is encounted with the merging of the .login file.

Number: Patch 754.00

Abstract: Fixes slow boot problems when booting cluster

State: Supersedes Patches 450.00, 752.00

- Ensures that a cluster member will be up to date with respect to the LSM configuration when calls are made to an internal LSM routine. Without this fix, smsd triggers LSM configuration errors when querying LSM in a cluster.
- Prevents LSM configuration daemon, vold, from core dumping when attempting to delete a disk that was not initialized properly.
- Fixes a very slow boot when booting several cluster nodes at the same time and CLSM is configured.

Number: Patch 756.00

Abstract: Correct improper file or privilege management

State: New

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file or privilege management.

Number: Patch 758.00

Abstract: Revision to the vdump.8 reference page

State: New

- Revises the vdump(8) reference page to change the statement for the -b option to be a maximum of 2048.

Number: Patch 760.00

Abstract: Corrects a problem in the marvel_pfm driver

State: New

- Corrects a problem in the marvel_pfm driver where xmesh and bmesh incorrectly reported 100 percent utilization for IO7 ports held in reset. Previously, this had to be corrected by using the mvfi test program.

Number: Patch 762.00

Abstract: Fix for SysMan Station failure

State: New

- Corrects a problem that occurred in some cluster configurations where SysMan Station fails to generate the hardware view. A Java stack trace is generated indicating that the routine "HardwareLayout.fancyPlace" was being executed at the time of the trace.

Number: Patch 764.00

Abstract: Correct improper file or privilege management

State: New

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file or privilege management.

Number: Patch 766.00

Abstract: Installs DECthreads V3.20-033

State: Supersedes Patches 22.00, 478.00

- Installs DECthreads V3.20-033 to address floating point problems in threaded programs. Previous patches installed DECthreads V3.20-029 and V3.20-029c.

Number: Patch 768.00

Abstract: Installs DECthreads V3.20-033

State: Supersedes Patches 24.00, 480.00

- Installs DECthreads V3.20-033 to address floating point problems in threaded programs. Previous patches installed DECthreads V3.20-029 and V3.20-029c.

Number: Patch 770.00

Abstract: Fixes cluster interconnect route advertisement

State: Supersedes Patch 42.00

- Corrects a problem using MD5 authentication with Version 2 RIP.
- Resolves a problem where the cluster interconnect route is inappropriately advertised.

Number: Patch 772.00

Abstract: Fix for find command

State: Supersedes Patch 64.00

- Corrects the find -ls command to display the correct number of blocks.
- Corrects the find -links, -size, -i, -inum behavior with respect to the + operations. Find + operations will match greater than, rather than greater than or equal to.

Number: Patch 774.00

Abstract: Fix for dirclean utility

State: New

- Corrects the /usr/sbin/dirclean utility from attempting to remove the AdvFS .tags directory or the quota.group and quota.user files.

Number: Patch 776.00

Abstract: Fixes spike code optimization tool

State: New

- Fixes the way the Spike code optimization tool handles -arch and -tune options. Using these options with previous versions of Spike resulted in unproductive error messages and/or system crashes.

Number: Patch 779.00

Abstract: Provides fixes for the collect utility

State: Supersedes Patches 195.00, 777.00

- Makes the following changes to the collect utility:
  - Updates the utility from Version 2.0.0 to 2.0.5.
  - Fixes several problems, including the handling Floating Point Exception.
  - Provides new collect features such as AdvFS monitoring and CPU and memory metrics on a per RAD basis.

Number: Patch 781.00

Abstract: Adds pmadvfs.mib to define AdvFS MIB definitions

State: New

- Adds a new MIB file, pmadvfs.mib, to define AdvFS MIB definitions.

Number: Patch 783.00

Abstract: Fix buffer overflow and improper file access

State: Supersedes Patches 246.00, 456.00

- Corrects several potential security vulnerabilities where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file or privilege management.

Number: Patch 785.00

Abstract: Fix for nissetup script

State: New

- Corrects a condition in which the nissetup command would leave /etc/group with an incorrect mode of 600 after removing NIS.

Number: Patch 787.00

Abstract: Revision of the fwupgrade(8) reference page

State: New

- Revises the fwupgrade(8) reference page to document new functions that allow the specified firmware update image to be located on a BOOTP server in a connected network.

Number: Patch 789.00

Abstract: Correct improper file or privilege

State: New

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file or privilege management.

Number: Patch 791.00

Abstract: Correct improper file access

State: Supersedes Patches 277.00, 464.00

- Corrects several potential security vulnerabilities where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file or privilege management.
- Corrects the behavior of the more command, when given both a nonexisting file and a nonempty file with a long file name.

Number: Patch 793.00

Abstract: Corrects I/O performance issue seen with CLSM

State: New

- Corrects an I/O performance issue that occurs when CLSM is configured and one member of a cluster goes down unexpectedly.

Number: Patch 795.00

Abstract: Fix for advfsstat command

State: New

- Corrects a problem with the advfsstat command printing negative values for statistics that are over 10 decimal digits long.

Number: Patch 797.00

Abstract: Fixes cut cmd to handle incomplete lines correctly

State: New

- Fixes /usr/bin/cut to correctly handle an incomplete line.

Number: Patch 801.00

Abstract: Fixes reset logic for Tru64 IDE/ATAPI driver

State: Supersedes Patches 440.00, 798.00, 799.00

- Fixes the problem of a kernel memory fault in systems that contain more than eight IDE/ATA buses.
- Fixes an IDE/ATA bus hang caused by attempting to complete raw odd byte DMA transfers to or from IDE/ATAPI devices.
- Prevents an IDE bus hang caused when issuing a play audio track command from scu to an ATAPI CD-ROM containing an enhanced CD.
- Fixes the IDE/ATAPI driver's reset logic to prevent a kernel memory fault when booting and to properly detect and log all master and slave reset failures when the system is operational.

Number: Patch 803.00

Abstract: Fix for migrate utility

State: New

- Corrects the error message returned when trying to migrate striped files when the -s option is omitted.

Number: Patch 805.00

Abstract: Sysman station display does not get updated

State: New

- Corrects a problem in which a Sysman Station display is not updated when component (for example, CPU) registration or deregistration events occur.

Number: Patch 809.00

Abstract: Fix for bcheckrc script

State: New

- Fixes a problem with bcheckrc that occurs when it is run multiple times.

Number: Patch 811.00

Abstract: Fix for defragment utility

State: New

- Fixes a problem where defragmentation can fail if the process can not obtain enough memory.

Number: Patch 813.00

Abstract: Fixes a problem in fixfdmn

State: Supersedes Patch 349.00

- Corrects several problems with the fixfdmn utility.
- Allows the fixfdmn utility to fix a rare corruption case in the RBMT/BMT0.

Number: Patch 815.00

Abstract: Addition of new sysconfigurable attribute

State: New

- Fixes a high lock contention for the str_to_lock STREAMS attribute.
- Adds a sysconfigurable tuning parameter to the STREAMS subsystem.

Number: Patch 817.00

Abstract: Correct improper file or privilege management

State: New

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file or privilege management.

Number: Patch 819.00

Abstract: Corrects problem with volmake utility

State: New

- Corrects a problem where, when creating a new plex, volmake will report that associating a subdisk with the plex would cause an overlap with another subdisk when they should not be overlapping.

Number: Patch 821.00

Abstract: Fixes a problem with scu

State: New

- Fixes a problem with the SCSI utility program, scu, where a mismatch between expected and found data displays incorrect data expected.

Number: Patch 823.00

Abstract: Revision to the sys_attrs_proc(5) reference page

State: New

- Revises the sys_attrs_proc(5) reference page to add the new tunable executable_data.
- Adds the new javaexecutedata(8) reference page .

Number: Patch 825.00

Abstract: Correct improper file access

State: Supersedes Patches 462.00, 294.00

- Provides support for SmartArray disk controllers. Without this patch, if the SmartArray product is installed on the system, the SysMan Station hardware view will fail to operate.
- Corrects a potential security vulnerability has where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file access.
- Provides enhancements for file system suitlets.

Number: Patch 835.00

Abstract: Correct improper file access

State: New

- Corrects a potential security vulnerability has where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file access.

Number: Patch 837.00

Abstract: Fix re_ioctl() cases DIODCMD and DIODCDB

State: New

- Fixes the re_ioctl() cases DIODCMD and DIODCDB where cmd transfer size has been changed to avoid kernel memory fault.

Number: Patch 838.00

Abstract:

State: Supersedes Patches 65.00, 66.00, 67.00, 68.00, 69.00, 70.00, 71.00, 72.00, 73.00, 75.00, 376.00, 406.00, 408.00, 488.00, 241.00, 242.00, 244.00, 435.00, 498.00, 213.00, 108.00, 109.00, 110.00, 111.00, 112.00, 113.00, 114.00, 115.00, 116.00, 117.00, 118.00, 119.00, 120.00, 121.00, 122.00, 123.00, 124.00, 125.00, 126.00, 127.00, 128.00, 129.00, 130.00, 131.00, 133.00, 397.00, 398.00, 399.00, 400.00, 401.00, 402.00, 403.00, 405.00, 490.00, 78.00, 79.00, 80.00, 82.00, 341.00, 409.00, 410.00, 412.00, 482.00, 494.00, 581.00, 582.00, 583.00, 584.00, 585.00, 586.00, 587.00, 588.00, 589.00, 590.00, 591.00, 592.00, 593.00, 594.00, 595.00, 596.00, 597.00, 598.00, 599.00, 600.00, 601.00, 602.00, 603.00, 604.00, 605.00, 606.00, 607.00, 608.00, 609.00, 610.00, 611.00, 612.00, 613.00, 614.00, 615.00, 616.00, 617.00, 618.00, 256.00, 25.00, 26.00, 27.00, 28.00, 30.00, 14.00, 161.00, 152.00, 153.00, 155.00, 389.00, 18.00, 507.00, 508.00, 509.00, 510.00, 511.00, 512.00, 513.00, 514.00, 515.00, 516.00, 517.00, 518.00, 519.00, 520.00, 521.00, 660.00, 523.00, 826.00, 827.00, 828.00, 829.00

- Provides more helpful informational messages for the following situations:
  - The advscan command will report if a domain has all of its volumes, but they are stored in a different directories — a situation that will cause the mount command mount to fail.
  - The AdvFS I/O error message will include the location of a file that will help users translate the error number into an error message.
- Fixes an rmvol E_PAGE_NOT_MAPPED error.
- Eliminates an ENO_MORE_BLKS error seen when COW'ing to a clone file while an rmvol operation is in progress.
- Fixes many small problems with the dsfmgr command.
- Provides protection against a class of potential security vulnerabilities called buffer overflows. Buffer overflows are sometimes exploited in an attempt to subvert the function of a privileged program and possibly execute commands at the elevated privileges if the program file has the setuid privilege. This patch allows a system administrator to enable memory management protections that limit potential buffer overflow vulnerabilities.

Number: Patch 839.00

Abstract: Fixes problems in dbx and object file tools

State: New

- Fixes various problems in the dbx and object file tools: dbx, ostrip, strip, mcs, dis, cord, file, and stdump.

Number: Patch 847.00

Abstract: Security (SSRT2266)

State: Supersedes Patches 179.00, 180.00, 181.00, 183.00, 427.00

- Fixes a potential security vulnerability that may result in denial of service. This may be in the form of local and remote security domain risks. The following potential security vulnerability has been corrected:

  SSRT2266 IGMP (Severity - High) which, under certain circumstances, could compromise system integrity.

- Fixes a problem in the kernel network subsystem that causes a kernel memory fault panic in the m_adj() routine.
- Fixes a problem in the IP multicast loopback code that causes a kernel memory fault panic.
- Fixes a problem in which the kernel incorrectly closes a socket, thereby causing Sybase 1613 errors to be produced.
- Fixes a problem in which a duplicate IP address might be configured on the system, or an IP address might be configured with an incorrect netmask.
- Fixes a problem caused when the Tru64 UNIX TCP layer prematurely closes a slow but good connection with TCP reset.

Number: Patch 851.00

Abstract: Fixes various problems in the bcm driver

State: Supersedes Patches 425.00, 43.00, 45.00, 48.00, 49.00, 50.00, 51.00, 52.00, 54.00, 500.00, 260.00, 567.00, 568.00, 569.00, 570.00, 571.00, 572.00, 574.00

- Adds IEEE 802.1Q Virtual Local Area Network (VLAN) support for the following:
  - DEGPA
  - DEGXA
  - DE50x, lan_common.h
  - DE60x
- Adds code to print greater than 61 UNIX domain sockets.
- Changes file read errors from /dev/kmem to ignore and continue in a running system.
- updates ddr.mod to support new hardware (NHD6) devices.
- Adds support to the ifconfig application for the IPv6 command line argument ip6reachabletime.
- Adds support for Ethernet adapters, including the DS25 onboard 10/100/1000 port.
- Fixes a problem in the alt driver for DEGPA Gigabit Ethernet adapters. This problem affects all Tru64 UNIX systems containing DEGPA network interfaces.
- Fixes numerous issues in the driver for DEGXA Gigabit Ethernet adapters.
- Fixes a bug that may cause a panic if the Xserver is stopped.
- Fixes a potential system crash when shutting down after using a DAPBA or DAPCA ATM adapter.
- Corrects a problem in which an ARP request for a permanent ARP entry is ignored and the user cannot connect from remote system.
- Corrects problems with netstat and ifconfig so that when a MAC address is printed, it is printed using 2-digit hex octets with leading zeros.
- Resolves a problem where some de50x network interface cards, under specific circumstances, may not send gratuitous ARP packets .
- Fixes a problem with non-U.S. USB keyboards used in non-U.S. locales in which the keyboards are treated as U.S. keyboards by the operating system.
- Fixes various problems in the bcm driver for DEGXA Gigabit Ethernet which can cause crashes.
- Adds recognition for possible future devices.

Number: Patch 859.00

Abstract: Fixes ftpd file transfer failure

State: New

- Fixes an ftpd file transfer failure when a file exeeds 2GB.

Number: Patch 861.00

Abstract: Fixes a problem in .mrg..protocols merge script

State: New

- Fixes a problem in the /etc/.mrg..protocols merge script that causes incorrect permissions on the /etc/protocols file.

Number: Patch 867.00

Abstract: Fixes a process hang problem in ubc_common_lookup

State: Supersedes Patches 290.00, 292.00, 436.00, 438.00, 92.00, (93.00, 94.00, 95.00, 96.00, 97.00, 98.00, 99.00, 100.00, 101.00, 374.00, 103.00, 417.00, 419.00, 496.00, 623.00, 624.00, 625.00, 626.00, 627.00, 628.00, 629.00, 630.00, 631.00, 632.00, 633.00, 635.00, 833.00, 845.00, 865.00

- Provides support for the SDLT160/320 and Ultrium 2 SCSI tape drives, including support for the Ultrium 2 SCSI to rewind after a reset behavior.
- Ensures proper compilation of the DDR database.
- Fixes a problem with the Smart Array driver that could cause a system hang to occur during error recovery when I/O is active.
- Adds support for new EVM events to be generated by the Event Monitoring daemon, /usr/sbin/envmond.
- Fixes the system panic "PWS_CCB_QUE_REMOVE: ccb not on any list," caused by a device or bus reset occurring during the execution of a command to a media changer device, like a tape library.
- Corrects a problem that causes a system panic while running applications performing open of a RAID device, and the faulting routine was control_port_open.
- Adds an event to indicate that the soft or hard error count has changed on the device identified in the event.
- Fixes a situation in which mounting a valid CD-ROM for the first time fails with the message “No valid file system exists on this partition,” although subsequent mounts of the same CD-ROM work as expected.
- Provides a configurable setting that causes an error return for any read of tape from a tape that requests less than the full amount of data in the tape block.
- Enables SmartArray 5300 controller hardware events to be logged to binary.errlog during a boot. This is useful in diagnosing logical volume state change and physical drive hotswaps that can occur while the system in not booted.
- Corrects a problem in which /sbin/ddr_config does not accept values for ReadyTimeSeconds larger than 255. The new limit is 86400 seconds (24 hours).
- Fixes problems with NUMA disk statistics.
- Fixes a KMF problem that can occur when some nodes in cluster are rebooted and a device is shared by all the nodes.
- Changes the CAM subsystem message that is printed to the error log on a recovered read error from "bad block number" to "block number."
- Corrects a problem in which camreport may report negative device IDs.
- Fixes the reporting to the binary errlog of device monitoring events and hardware errors during disk recovery from the disk driver.
- Corrects a problem with hwmgr utility deletes while a SCSI scan is in progress.
- Corrects a problem in which a path event can cause hang in cdisk_online during disk open of HSG80.
- Installs the V1.07 release of the ciss driver, which is the mandatory minimum version to support the Smart Array 5300 Controller.
- Addresses an issue where AdvFS domain panics would occur during HSZ and HSG failovers.

Patch 867.00 Continued

- Fixes a problem where the CAM I/O subsystem does not always zero the Cam Control Blocks that are used by the peripheral drivers. This can cause a kernel memory fault or system hang when the subsystem is low on memory.
- Prevents unnecessary retries on an HSG80 when fail unit attention with ascq = Oxf002 and returns poper error to higher layer.
- Prevents "ccfg_MakeDeviceIdentWWID: Invalid device ID" messages from being generated when they should not be.
- Provides Version 1.08 of the ciss driver.
- Verifies path structures in ctape_ioctl and ctape_generic_passthru to prevent a kernel memory fault if the tape was opened with FNDELAY flag set.
- Addresses an issue where AdvFS domain panics would occur during HSZ and HSG failovers.
- Fixes a small memory leak in Power Management code.
- Provides Version 1.09 of the ciss driver.
- Corrects a problem in which the BBR code logs all error messages as soft error, even if the error was not recovered and it failed to do the bad block replacement.
- Fixes a process hang in ubc_common_lookup.
- Fixes a problem where SmartArray 5300 Logical Volumes were counted as RAID controllers.
- Addresses an issue where NULL Inquiry data causes a "'Device has no 'name'" error as well as possible I/O stalls.
- Provides Version 1.10 of the ciss driver.
- Corrects problems where tape devices become unavailable, do not respond, report unrecoverable errors, or cause kernel memory faults.
- Adds version 1.12 of the ciss driver.
- Fixes a problem in which a GS80 would hang following an SA5300 reset.
- Addresses an issue in which I/O may not complete under certain circumstances.

Number: Patch 872.00

Abstract: Security (SSRT2301, SSRT2275, SSRT2384)

State: Supersedes Patches 65.00, 66.00, 67.00, 68.00, 69.00, 70.00, 71.00, 72.00, 73.00, 75.00, 376.00, 406.00, 408.00, 488.00, 241.00, 242.00, 244.00, 435.00, 498.00, 213.00, 108.00, 109.00, 110.00, 111.00, 112.00, 113.00, 114.00, 115.00, 116.00, 117.00, 118.00, 119.00, 120.00, 121.00, 122.00, 123.00, 124.00, 125.00, 126.00, 127.00, 128.00, 129.00, 130.00, 131.00, 133.00, 397.00, 398.00, 399.00, 400.00, 401.00, 402.00, 403.00, 405.00, 490.00, 78.00, 79.00, 80.00, 82.00, 341.00, 409.00, 410.00, 412.00, 482.00, 494.00, 581.00, 582.00, 583.00, 584.00, 585.00, 586.00, 587.00, 588.00, 589.00, 590.00, 591.00, 592.00, 593.00, 594.00, 595.00, 596.00, 597.00, 598.00, 599.00, 600.00, 601.00, 602.00, 603.00, 604.00, 605.00, 606.00, 607.00, 608.00, 609.00, 610.00, 611.00, 612.00, 613.00, 614.00, 615.00, 616.00, 617.00, 618.00, 256.00, 25.00, 26.00, 27.00, 28.00, 30.00, 14.00, 161.00, 152.00, 153.00, 155.00, 389.00, 18.00, 507.00, 508.00, 509.00, 510.00, 511.00, 512.00, 513.00, 514.00, 515.00, 516.00, 517.00, 518.00, 519.00, 520.00, 521.00, 660.00, 523.00, 826.00, 827.00, 828.00, 829.00, 831.00, 840.00, 841.00, 843.00, 849.99, 862.00, 864.00, 869.00, 870.00

- Fixes a process hang condition.
- Fixes a "thread_block: simple lock held" panic.
- Corrects a situation in which a system could panic during a particular machine check.
- Corrects several problems of 3D client hangs when using a Radeon graphics card.
- Fixes a performance problem seen when doing wiring on gh_chunks memory; for example, an Oracle application.
- Protects against "get_color_bucket: empty buckets!" panics and "kernel memory fault" failures on systems with mixed cache parameters.
- Fixes a kernel memory fault in u_seg_global_destroy.
- Corrects a kernel memory fault that can happen when running applications that use the Cray Intra-Node Shared Memory library.
- Prevents a potential process (not system) hang seen when a system comes under heavy memory load with monolithic memory use (gigabyte-scale single objects).
- Prevents a kernel memory fault when running with protection on the 128-byte bucket. (This should only be running with this as directed by support personnel.)
- Corrects a situation in which a taso-compiled binary is unable to allocate more memory after performing a series of mmap calls.
- Fixes an occasional panic that can be seen when reading from a process using Granularity Hints via the /proc file system.
- Fixes a panic that generates the message “u_seg_vop_remove: seg not found in vop.”
- Fixes a situation in which mmap memory locked with mlockall() using the MCL_FUTURE flag does not become wired automatically.
- Fixes a "Bigpage Assertion Failed" panic.
- Corrects a rounding error for vm attribute vm_bigpg_thresh.
- Corrects the handling of bad pages when bigpages are enabled.
- Fixes "page mapped" panics when using the mmap() function for dev/mem to access free bigpages.

Patch 872.00 Continued

- Corrects a condition that causes a delete_pv_entry panic when kernel virtual-address space has high usage.
- Fixes a problem seen when USB hubs (or any other bus device) are removed from a running system.
- Removes a restriction in which dynamic VMEbus device drivers could only probe one controller per driver. With this patch, multiple controllers per driver can be configured successfully.
- Fixes a potential floating point register corruption.
- Fixes multiple problems affecting a system with peripheral USB hubs attached, as well as problems that might occur when moving or adding USB host adapters.
- Improves I/O performance by reducing kernel locking overhead.
- Fixes a system hang when using Open3D over the AGP bus on a GS1280.
- Corrects performance issues when accessing a file with direct I/O enabled.
- Fixes a panic that can occur when appending to a file.
- Fixes an AdvFS asynchronous direct I/O problem that could cause a thread to hang.
- Fixes a problem encountered where a truncated AdvFS file erroneously zeroed data for the remaining leading segment of the file.
- Changes the behavior of migrate_normal and migrate_stripe when migrating an original file that has a clone. If the clone was marked out of sync, migrate could come back with E_CLONE_OUT_OF_SYNC even though the migrate succeeded. Now this case is caught, and handled.
- Replaces the system panics caused by "Can't clear bit twice" with a domain panic.
- Fixes a problem in which a crash occurs when an AdvFS file system reports I/O errors and enters into a domain panic state. AdvFS's error cleanup would panic on an invalid pointer and report an "invalid memory read access from kernel mode" panic message.
- Fixes an issue encountered in configurations where the primary processor is not the first processor within a rad.
- Resolves a problem of not being able to view files on some CD-ROM media that is created by third-party software and corrects the erroneous reporting of success when attempting to write beyond the file size limit using synchronized I/O and the calculation of _PC_FILESIZEBITS, which is used by the operating system for pathconf file characteristics.
- Fixes a problem with audit data not being displayed by the audit tool,
- Fixes problems with file object selection and deselection and directories.
- Fixes problems with NUMA performance associated with auditing.
- Fixes a race during AdvFS volume removal that can cause a panic in the bs_osf_complete() routine.
- Fixes a problem with kernel memory fault in shadowvnode() caused by NULL vnode pointer.
- Fixes insmntque() to conform to proper locking when removing and adding vnode to the mount vlist.
- Fixes a problem with excessive FIDS_LOCK contention observed when large numbers of files are using system-based file locking.

Patch 872.00 Continued

- Corrects the AdvFS system call OP_GET_BKUP_XTNT_MA to avoid a silent infinite loop in vdump. The call will now return the valid xtntCnt when it fails due to E_NOT_ENOUGH_XTNTS.
- Fixes a panic caused by a problem within the swapping subsystem.
- Fixes mount and umount failures and panics in MFS, UFS, and FDFS.
- Fixes an AdvFS alignment fault panic caused by inconsistent AdvFS metadata in a directory. In particular, the directory's entry size is an unaligned value.
- Corrects a potential problem with modifying files via direct I/O when there is a clone fileset.
- Fixes a panic within the two-level scheduling subsystem.
- Improves AdvFS informational messages as follows:
  - Advscan reports if a domain has all of its volumes, but they are stored in a different directories. This scenario will cause mount to fail.
  - The AdvFS I/O error message includes the location of a file that will help users to translate the error number into an error message.
- Forces a domain panic instead of a system panic if AdvFS metadata is discovered to be incorrect in frag_group_dealloc.
- Fixes a part of AdvFS migration code in order to prevent the rmvol utility's "Can't remove volume" error.
- Fixes a problem when monitoring I/O using the advfsstat command.
- Fixes a rare problem that causes thread blocking when waiting for memory.
- Fixes a problem in which a domain panic may occur in idx_lookup_node_int or bs_frag_dealloc under heavy file system activity, generating one of the following messages:

  idx_lookup_node_int: bs_refpg failed
  bs_frag_dealloc: rbf_pinpg (4) failed, return code = -1035

- Fixes a problem in which the fuser utility is unable to report on all referenced resources, which occurs when attempting to identify reasons for unmount failures.
- Improves the process exit procedure for processes that have had the nice command used on them.
- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file access.
- Corrects a race condition in AdvFS in which it avoids a potential stranded log record in memory that does not get out to disk.
- Fixes an rmvol E_PAGE_NOT_MAPPED error.
- Eliminates an ENO_MORE_BLKS error seen when performing a copy on write (COW) procedure to a clone file while an rmvol operation is in progress.
- Fixes a problem in which a system on a cluster can panic with the message "ics_unable_to_make_progress: input thread stalled."
- Adds support for CPU indictment on AlphaServer ES80 and GS1280 systems.
- Adds support in the platform code to handle MSI capable adapters. AlphaServer GS1280 systems support option cards that require MSI capabilities.

Patch 872.00 Continued

- Adds support to get live status information for air movers and power supplies on AlphaServer ES80 and GS1280 systems and to log intrusion packets to the error log.
- Fixes a problem in which a process waiting on a semaphore does not get woken up.
- Fixes a problem in which the extension of UNIX file systems via the mount command can effectively disable the use of the file system.
- Fixes a problem on some LSM-based systems in which a panic can occurs after a file system extension has been completed.
- Fixes a problem in which a hang may occur when a rmvol operation is performed after a cluster node failure during volmigrate, volunmigrate, or frag file migration.
- Corrects a locking problem with NFS running over UFS.
- Fixes an obliteration of user file information, which is most often seen after using the ftruncate() function.
- Fixes two potential problems in the NFS V3 client where unstable writes could potentially remain uncommitted when they should have been committed to stable storage.
- Eliminates a false directory lookup warning message generated by an incorrect comparison caused by mismatched fileid variable types.
- Improves client caching performance.
- Fixes a problem that can cause a system crash when an NFS server exports files on a third-party file system (that is, one not built into Tru64 UNIX).
- Prevents the loss of a single system image for an NFS file system mounted from a cluster, as a result of problems communicating with the external NFS server.
- Fixes a memory leak in the NFS server when it receives malformed packets.
- Allows the size of the NFS server's duplicate request cache to be adjusted as needed.
- Fixes a problem in which a Tru64 UNIX NFS client panics when it receives a null entry as a response to a readdirplus request from an NFS server.
- Fixes a problem in which a Tru64 UNIX NFS client panics as a result of receiving illegal file access mode from an NFS client.
- Increases TCP windows from 96 KB to 500 KB to improve performance.
- Causes the netisr thread to dynamically estimate the reply size and subsequently reserve the space in the socket buffer.
- Adds a new timeout check to notice when data in a socket buffer has not been acknowledged in 30-50 seconds and copies those buffers to allow the UBC to free up those mbufs and not tie them up.
- Fixes a flaw in the NFS server that could cause it to crash upon reception of malformed input.
- Addresses the following problems with the NFS server:
  - A potential crash with concurrent read and truncate on an AdvFS file
  - A potential crash with malformed or malicious READDIR[PLUS] version 3 RPCs
- Corrects problems with the time of year (TOY) clock.
- Prevents a panic that may occur in a cluster when a fileset mounted -o dual is failed over or unmounted during shutdown.
- Ensures that time is kept accurately when a system is configured to use NTP, regardless of whether the NTP daemon is running.
- Prevents a recursive panic situation on a ES47 platform when a double bit memory error is detected.

Patch 872.00 Continued

- Corrects the problem of a possible panic in audit_rec_build when auditing execve with the exec_argp or exec_envp audit style enabled.
- Increases the the maximum length of names in a file property lists from 245 to 255 characters.
- Fixes a problem where gh_min_seg_size could not be set below 8M.
- Improve the performance of AdvFS under heavy I/O loads.
- Allows CPUs that do not take interrupts (either directly from the attached IO7 or indirectly from IO7s attached to offlined CPU) to be taken off line.
- Prevents issues in the DCE/DFS file system when pages are being flushed as part of a vnode. This patch is required for Sierra systems.
- Allows AdvFS to record if a domain panic has occurred, even if a system panic results.
- Helps prevent kernel memory faults in AdvFS.
- Fixes an AdvFS path that can cause a panic in the advfs_page_busy() routine.
- Fixes a hang that can occur during the renaming of an AdvFS file.
- Fixes a system panic in the ubc_page_stealer routine.
- Fixes a problem that causes a panic that results from a race condition in MFS (memory file system) over CFS (cluster file system).
- Fixes a problem with cluster failover of UFS filesets.
- Prevents a O_DSYNC write failure under the following situation:

  1. The user creates a new file
  2. Closes the file
  3. The vnode for the file is recycled
  4. Reopens file with O_DSYNC flag
  5. Writes to file, overwriting already allocated storage
  6. The write from step 5 returns to the application.
  7. The system crashes

- Replaces two potential panics in AdvFS with domain panics.
- Changes the fwupgrade command to allow the specified firmware update image to be located on a BOOTP server in a connected network.
- Provides protection against a class of potential security vulnerabilities called buffer overflows. Buffer overflows are sometimes exploited in an attempt to subvert the function of a privileged program and possibly execute commands at the elevated privileges if the program file has the setuid privilege. This patch allows a system administrator to enable memory management protections that limit potential buffer overflow vulnerabilities.dds only comment to reserve 0x10000000 and 0x20000000 for AUTOFS flags
- Adds comment to reserve 0x10000000 and 0x20000000 for AUTOFS flags.
- Corrects an internal AdvFS check that was always returning true.
- Fixes a deadlocking problem in the kernel where a file open on a clone could hang when ACLs are enabled.
- Includes some scalability improvements to AdvFS that will help reduce lock contention and improve performance.

Patch 872.00 Continued

- Prevents a potential hang during reboot if there had been a recent domain panic.
- Addresses small memory leaks within the kernel that occur infrequently.
- Sisplays the correct error message for the freezefs -q command on a non-AdvFS file system.
- Addresses system problems that can occur when the system is under heavy I/O load and/or low memory conditions.
- Fixes a standard violation on AdvFS.
- Corrects an a "ialloc: dup alloc" panic.
- Corrects a "blkfree: freeing free block" panic and a "blkfree: freeing free frag" panic.
- Improves the scalabililty and performance of AdvFS.
- Fixes a problem where an I/O error (EIO) can occasionally be returned after a page fault.
- Fixes a possible kernel system hang in vfast when shutting down or rebooting the system.
- Prevents a kernel memory fault panic that would occur when the audit daemon is set to periodically dump the kernel audit buffers to the audit log file (auditd -d freq).
- Adds defensive programming to stat.h to avoid stat.h getting confused if one of its internal temporary #defines is defined before stat.h is processed.
- Fixes a memory management fault/panic in UFS.
- Fixes a problem where a device file such as /dev/console can become inaccessible, returning the error “Bad File Number.”
- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file or privilege management.
- Allows the auditing of login and su events based in part on the contents of user profiles (for Enhanced Security), the prevailing auditing characteristics of the originating process, and the system-wide audit mask. Previously, only the system audit mask was referenced.
- Corrects a failure in the safe_open() routine that caused symbolic links given by a relative path from the current working directory sometimes to give ENOENT errors incorrectly.
- Fixes a potential floating point error in threaded applications.
- Fixes an extended regular expression problem where the interval expression {m,n} is handled incorrectly.
- Fixes a problem with SIA that caused the Internet Express LDAP Authentication module to be unable to look up default group information for a user at login time.
- Fixes many small problems with the dsfmgr command.
- Corrects a cp command performance problem.
- Corrects a problem in which the cp and cat commands produce different file sizes when reading from a tape device.
- Corrects a potential security vulnerability has where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file access.
- Corrects a problem in which the sh shell was using a high amount of CPU time.
- Fixes the problem while encoding $@ in the Bourne shell.
- Addresses several problems with ssh secure shell, including interoperability with other ssh implementations, ssh client/server configuration files compatibility issues, and the lack of IPv6 support.

Patch 872.00 Continued

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file or privilege management.
- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised when a buffer overflow occurs in the uucp utility. Buffer overflows are sometimes exploited in an attempt to subvert the function of a privileged program and possibly execute commands at the elevated privileges if the program file has the setuid privilege.
- Provides protection against a class of potential security vulnerabilities called buffer overflows. Buffer overflows are sometimes exploited in an attempt to subvert the function of a privileged program and possibly execute commands at the elevated privileges if the program file has the setuid privilege. This patch allows a system administrator to enable memory management protections that limit potential buffer overflow vulnerabilities.
- Allows the sh shell to print the correct msg when enhanced core file naming is on.
- Fixes a problem in which the runtime loader's (/sbin/loader) attempts to free a null pointer are in error.
- Corrects RPC-based servers' handling of ill-formed TCP connections.
- Corrects a problem in which the return value of unlink() call was not checked when two threads were trying to move a file to two different destinations. Although one of the threads could unlink() the source file, no relevant error message was displayed.
- Addresses a potential security vulnerability that may may result in a Denial of Service (DoS). This potential vulnerability may be in the form of local and remote security domain risks. The following potential security vulnerability has been corrected: SSRT2384 rpc (Severity - High)
- Allows mount(8) options that take a value to be correctly processed on a cluster.
- Fixes memory leaks caused by certain type of scripts that called an infinite loop.
- Fixes following problems in sh:
  - Service denial problem when a quoted here doc script is executed.
  - Problem with handling ELF files.
  - The shell variable $- not holding -C set option when it is turned on.
  - Printing broken characters when type builtin utility of sh is invoked in Japanese locale.
- Prevents segmentation faults when sia_ses_init is passed a malformed argument vector.
- Fixes client (login, su, rshd, edauth, and sshd2) hangs and long delays under Enhanced Security, as well as some intermittent errors or failures seen with prpasswdd or rpc.yppasswdd.
- Fixes various problems in the libc functions getdate(), strptime(), callrpc(), strncasecmp(), and fork(). It also fixes a problem in the libnuma function nacreate() and the system header <sgtty.h>.
- Fixes a problem where the home directory and login shell attributes for a user account were not suppled to the audit daemon for authentication failures.
- Fixes various problems in the dbx and object file tools: dbx, ostrip, strip, mcs, dis, cord, file, and stdump.
- Corrects a problem in which some networking applications, especially X.25 and X.29, stopped working as expected because of interactions with security-related fixes and how the fstat() function behaves on their sockets.

Patch 872.00 Continued

- Fixes a regression from pre-V5.0 releases in the libc mktime() function's handling of potentially ambiguous tm struct times; that is, those that fall within a backward clock shift and that have an initially negative tm_isdst value.
- Fixes an internal problem in the kernel's AdvFS, UFS, and NFS file systems where extended attributes with extremely long names (greater than 247 characters) could not be set on files. The new limit is 254 + a Null string terminator.
- Corrects the creation of console boot device strings for devices on subordinate buses.
- Fixes the garbage character that sometimes appears when requesting the name of a boot device via consvar.
- Corrects a potential security vulnerability which may result in nonprivileged users gaining unauthorized access to files or privileged access on the system. This potential vulnerability may be in the form of a local and remote security domain risk.
- Fixes a sh shell problem with command substitution.
- Modifies console callback code to allow users to use upper and lower case variable names for known console environment variables. This patch is required for update installations on EV7 based platforms.
- Fixes a system panic that has the panic string "pg_nwriters going negative."
- Fixes a system panic. The panic string could be "mcs_lock: lock already owned by cpu" or "thread_block: simple lock owned."
- Fixes a problem on a cluster NFS client where a hard-mounted NFS file system can incur ETIMEDOUT errors.
- Corrects problems in UFS extendfs functionality, which could cause file system metadata inconsistency.
- Addresses an issue on large systems where kernel threads might not be executed for extended periods of time.
- Corrects a kernel memory fault in the table syscall.
- Resolves a problem that results in multiple cluster members crashing with a kernel memory fault in rfs_find_fsid().
- Reduces the delay with hwmgr of EV7 based machines by reducing the number of calls to the console to update sensor data.
- Fixes a problem in which the quot -v command sometimes returns wrong quota information on a UFS partition.

Number: Patch 874.00

Abstract: Security (SSRT3469)

State: New (Supercedes Patch 857.00)

- Corrects a security vulnerability in sendmail that might result in nonprivileged users gaining unauthorized access to files or privleged access. This potential vulnerability may be in the form of a local or remote security domain risk.

Number: Patch 876.00

Abstract: Revises collect(8) and pmgrd(8) reference pages

State: New

- Revises collect(8) and pmgrd(8) reference pages. See Section 1.1.9.

2

TruCluster Server Patches

This chapter provides information about the patches included in Patch Kit 2 for the TruCluster Server software.

This chapter is organized as follows:

- Section 2.1 provides release notes that are specific to the TruCluster Server software patches in this kit.
- Section 2.2 provides brief descriptions of the TruCluster Server patches included in this kit.

2.1 Release Notes

This section provides release notes that are specific to the TruCluster Server software patches in this kit.

2.1.1 Required Storage Space

The following storage space is required to install the base and TruCluster Server components of this patch kit:

- Approximately 250 MB of temporary storage space is required to untar this patch kit (base and TruCluster). We recommend that this kit not be placed in the `/`, `/usr`, or `/var` file systems because doing so may unduly constrain the available storage space for the patching activity.
- Approximately 89.3 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.
- Approximately 90.7 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.
- Approximately 1224 KB of storage space is required in `/var/adm/patch/doc` for patch abstract and README documentation.
- Approximately 176 KB of storage space is needed in `/usr/sbin/dupatch` for the patch management utility.

See Section 1.1.1 for information on space needed for the operating system patches.

2.1.2 AlphaServer ES47 or AlphaServer GS1280 Hangs When Added to Cluster

If after running `clu_add_member` to add an AlphaServer ES47 or AlphaServer GS1280 as a member of a TruCluster the AlphaServer hangs during its first boot, try rebooting it with the original V5.1B generic cluster kernel, `clu_genvmunix`.

Use the following instructions to extract and copy the V5.1B cluster `genvmunix` from your original Tru64 UNIX kit to your AlphaServer ES47 or AlphaServer GS1280 system. In these instructions, the AlphaServer ES47 or AlphaServer

GS1280 is designated as member 5. Substitute the appropriate member number for your cluster.

1. Insert the Tru64 UNIX Associated Products Disk 2 into the CD-ROM drive of an active member.
2. Mount the CD-ROM to `/mnt`. For example:

```
# mount -r /dev/disk/cdrom0c /mnt
```

3. Mount the bootdisk of the AlphaServer ES47 or AlphaServer GS1280 on its specific mount point; for example:

```
# mount root5_domain#root /cluster/members/member5/boot_partition
```

4. Extract the original `clu_genvmunix` from the CD-ROM and copy it to the bootdisk of the AlphaServer ES47 or AlphaServer GS1280 member.

```
# zcat < TCRBASE540 | ( cd /cluster/admin/tmp; tar -xf -
./usr/opt/TruCluster/clu_genvmunix)
# cp /cluster/admin/tmp/usr/opt/TruCluster/clu_genvmunix \
/cluster/members/member?/boot_partition/genvmunix
# rm /cluster/admin/tmp/usr/opt/TruCluster/clu_genvmunix
```

5. Unmount the CD-ROM and the bootdisk:

```
# umount /mnt
# umount /cluster/members/member5/boot_partition
```

6. Reboot the AlphaServer ES47 or AlphaServer GS1280.

2.1.3 No-Roll Procedure Cannot Be Used to Remove Patch Kit

To remove Patch Kit 2, you must run the `/etc/dn_fix_dat.sh` script after rebuilding the kernel and before rebooting the system. If the script is not executed before rebooting, the system will fail to boot.

Because the no-roll procedure automatically reboots the system after deleting the patches, you would not be able to run the script as required. Therefore, the no-roll procedure cannot be used to remove the patch kit.

The workaround is to use the rolling upgrade procedure to remove the patch kit. See Section 1.1.3 for more information.

2.1.4 Updates for Rolling Upgrade Procedures

The following sections provide information on rolling upgrade procedures.

2.1.4.1 Noncritical Errors

During a rolling upgrade to installing Patch Kit 2, you may encounter the following noncritical situations:

- The tagged file for `ifaccess.conf` (`.Old..ifaccess.conf`) may disappear. This error will not cause any problems with the rolling upgrade procedure or the installation of the kit. A message would alert you to this condition if you use the `clu_upgrade undo` command. Running the `clu_upgrade -v check setup` at the start of the procedure will fix this error.
- When the world-wide language subset is installed, the file `wwinstall` will attempt to be tagged, and will fail. This error will not affect the operational status of the cluster.

2.1.4.2 Procedure for Simultaneous Upgrades

When doing a simultaneous rolling upgrade of NHD6 and the Version 5.1B Patch Kit 2, you must install the NHD6 kit first. If this is not done, you may get a number of installation errors, although you can safely ignore them.

2.1.4.3 Unrecoverable Failure Procedure

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

2.1.4.4 Do Not Add or Delete OSF, TCR, IOS, or OSH Subsets During Roll

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

- Base Operating System subsets (those with the prefix `OSF`).
- TruCluster Server subsets (those with the prefix `TCR`).
- Worldwide Language Support (WLS) subsets (those with the prefix `IOS`).
- New Hardware Delivery (NHD) subsets (those with the prefix `OSH`).

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

2.1.4.5 Undo Stages in Correct Order

If you need to undo the install stage, because the lead member is in an unrecoverable state, be sure to undo the stages in the correct order.

During the install stage, `clu_upgrade` cannot tell whether the roll is going forward or backward. This ambiguity incorrectly allows the `clu_upgrade undo preinstall` stage to be run before `clu_upgrade undo install`. Refer to the *Patch Kit Installation Instructions* for additional information on undoing a rolling patch.

2.1.4.6 Ignore Message About Missing ladebug.cat File

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

```
Creating tagged files.

.....................................................................
.....

*** Error ***
The tar commands used to create tagged files in the '/usr' file system have
reported the following errors and warnings:
     tar: lib/nls/msg/en_US.88591/ladebug.cat : No such file or directory
.....................................................

*** Warning ***
The above errors were detected during the cluster upgrade. If you believe that
the errors are not critical to system operation, you can choose to continue.
If you are unsure, you should check the cluster upgrade log and refer
to clu_upgrade(8) before continuing with the upgrade.
```

2.1.4.7 clu_upgrade undo of Install Stage Can Result in Incorrect File Permissions

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

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

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

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

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

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

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

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

2.1.4.8 Missing Entry Messages Can Be Ignored During Rolling Patch

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

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

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

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

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

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

2.1.4.9 Relocating AutoFS During a Rolling Upgrade on a Cluster

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

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

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

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

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

```
named          application    ONLINE    ONLINE    rye
```

To minimize your effort in the following procedure, perform the roll stage last on the member where AutoFS runs.

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

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

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

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

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

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

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

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

2. Continue with the rolling upgrade procedure.

2.1.5 Additional Steps Required When Installing Patches Before Cluster Creation

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

1. Install the Tru64 UNIX base kit.
2. Install the TruCluster Server kit.
3. Install the Version 5.1B 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**

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

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

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

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

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

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

2.1.7 Problems with clu_upgrade switch Stage

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

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

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

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

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

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

2.2 Summary of TruCluster Software Patches

This section provides brief descriptions of the patches in Patch Kit 1 for the software products.

This section provides descriptions of the patches in Patch Kit 2 for the TruCluster Server software products. Because Tru64 UNIX patch kits are cumulative, each patch lists its state according to the following criteria:

- New

  Indicates a patch that is new for this release.

- New (Supersedes Patches ...)

  Indicates a patch that is new to the kit but was combined (merged) with one or more patches during the creation of earlier versions of this kit, before it was publicly released.

- Existing (Kit 1)

  Indicates a patch that was new in the previous Version 5.1B patch kit.

- Supersedes Patches ...

  Indicates a patch that was combined (merged) with other patches.

Number: Patch 2.00

Abstract: Fix for aliasd daemon

State: Existing (Kit 1)

Modifies the aliasd daemon to include interface aliases when determining whether or not an interface is appropriate for use as the ARP address for a cluster alias when selecting the proxy ARP master.

Number: Patch 7.00

Abstract: Fixes an issue with ICS on NUMA-based systems

State: Existing (Kit 1)

- Fixes an issue with ICS (Internode Communication Services) on a NUMA-based system in a cluster.

Number: Patch 14.00

Abstract: Cluster specific fix for mounting cluster root domain

State: Existing (Kit 1)

This patch enables a cluster to boot even if the cluster root domain devices are private to different cluster members. Although this is not a recommended configuration, it should not result in an unbootable cluster. Currently, this is with respect to cluster root domains not under LSM control.

Number: Patch 19.00

Abstract: Fix for Oracle startup failure

State: Existing (Kit 1)

- Fixes a problem in one of the shipped rc scripts whereby Oracle fails during startup on a clustered system.

Number: Patch 26.00

Abstract: Problems with LSM disks and cluster quorum tool

State: Existing (Kit 1)

- Corrects problems with LSM disks and the cluster quorum tools. When a member having LSM disks local to it is down, the quorum tools fail to update quorum. This causes other cluster commands to fail.

Number: Patch 35.00

Abstract: Fix for cluster alias manager SUITlet

State: Existing (Kit 1)

- Fixes a problem that causes the cluster alias manager SUITlet to falsely interpret any cluster alias with virtual={t|f} configured as a virtual alias regardless of its actual setting.

Number: Patch 39.00

Abstract: Reliable DataGram kernel thread problem

State: Existing (Kit 1)

- Fixes a problem in which an RDG (Reliable DataGram) kernel thread can starve other timeshare threads on a uniprocessor cluster member. In particular, system services such as networking threads can be affected.

Number: Patch 52.00

Abstract: Patch: Fix for RM_AUDIT_ACK_BLOCK

State: Supersedes Patches 3.00, 5.00

- Fixes a regression for single physical rail Memory Channel configurations, and cleans up stale data left on an offline physical rail by the Memory Channel driver.
- Fixes issues associated with the initialization of the Memory Channel driver.
- Corrects a problem in a Memory Channel cluster where rebooting a node without performing a hardware reset can crash other members with a RM_AUDIT_ACK_BLOCK panic.

Number: Patch 63.00

Abstract: Cluster member panics with Kernel Memory Fault

State: Supersedes Patches 15.00, 17.00

- Fixes a problem in which cluster alias connections are not distributed among cluster members according to the defined selection weight.
- Fixes a memory leak in the cluster alias subsystem.
- Corrects a problem that occurs when running nmap or nessus targetted at the cluster alias, where the cluster member panics with Kernel Memory Fault.

Number: Patch 65.00

Abstract: Resolves problem with caa_register command

State: New

- Resolves a problem in which the caa_register command allowed a CAA resource to be registered even when its profile contained an unknown attribute. This fix prevents the caa_register command from registering a resource with an unknown attribute and will cause it to return an error message which includes the unknown attribute information.

Number: Patch 67.00

Abstract: Fixes a cfsd core dumping problem

State: Supersedes Patch 48.00

- Fixes a problem with cfsd core dumping shortly after startup if it is enabled or shortly after enabling it. The problem fixed by this patch is only seen after applying a recent dsfmgr patch.
- Corrects a problem in which cfsd will terminate prematurely and core dump when a node leaves the cluster very shortly after joining the cluster.

Number: Patch 69.00

Abstract: Fixes cluster interconnect

State: Supersedes Patches 20.00, 22.00

- Corrects a problem involving discarded UDP datagrams that do not come from the correct port.
- Corrects a problem in which a panic displaying the message “error CNX MGR: cnx_comm_error: invalid node state” occurs on a LAN cluster running under load when other members are rebooting.
- Fixes a coding error, a memory leak, and a deinitialization problem in the cluster interconnect networking layer.

Number: Patch 72.00

Abstract: Fixes race condition in Device Request Dispatcher

State: Supersedes Patches 27.00, 28.00, 29.00, 31.00, 50.00, 70.00

- Fixes a regression associated with non-SCSI storage.
- Improves the responsiveness of EINPROGRESS handling during the issuing of I/O barriers by removing a possible infinite loop scenario that could occur due to the deletion of a storage device.
- Fixes a problem that causes a panic with the message "CNX MGR: Invalid configuration for cluster seq disk" during simultaneous booting of cluster nodes.
- Fixes a possible race condition between a SCSI reservation conflict and an I/O drain, which could result in a hang.
- Alleviates a condition in which a cluster member takes an extremely long time to boot when using LSM.
- Fixes a problem in the cluster kernel where a cluster member panics while doing remote I/O over the interconnect.
- Corrects an issue to allow the Device Request Dispatcher, DRD, to retry to get disk attributes when EINPROGRESS is returned from the disk driver.
- Fixes an problem where access to the quorum disk can be lost if the quorum disk is on a parallel SCSI bus and multiple bus resets are encountered.
- Fixes several problems in the Device Request Dispatcher, including a race condition.

Number: Patch 74.00

Abstract: Fix for caa_report

State: New

- Fixes a condition in which uptimes greater than 100 percent are reported for resources by caa_report.
- Fixes a problem in which resources that never started have an ending timestamp.

Number: Patch 81.00

Abstract: Security (SSRT2265)

State: Supersedes Patch 37.00

- Fixes a security vulnerability in the cluster interconnect security configuration that may result in a denial of service on systems running TruCluster Server software.
- Provides enhancements to the clu_upgrade command.

Number: Patch 85.00

Abstract: Fixes a panic that may occur during an unmount

State: Supersedes Patches 8.00, 9.00, 10.00, 12.00, 41.00, 44.00, 46.00, 53.00, 54.00, 55.00, 56.00, 57.00, 58.00, 59.00, 61.00, 83.00

- Fixes a problem that causes a hang to occur when multiple nodes are shutting down simultaneously.
- Fixes a problem that causes a Cluster File System panic when using raw Asynchronous I/O.
- Adds code to assist in problem diagnosis.
- Relieves pressure on the CMS global DLM lock by allowing AutoFS auto-unmounts to back off.
- Updates the attributes on a directory when files are removed by a cluster node that is not the file system server.
- Fixes a problem of excessive FIDS_LOCK contention that occurs when large number of files are using system-based file locking.
- Fixes a cluster deadlock that may occur during failover and recovery when direct I/O is in use.
- Corrects diagnostic code that might result in a panic during kernel boot.
- Prevents a panic when an AutoFS file system is auto-unmounted.
- Enhances cluster file system performance when using file locks to coordinate file access.
- Corrects several problems with various installation commands and utilities.
- Fixes a memory leak in the clu_get_info interface.
- Displays the correct error message for freezefs -q on a non-AdvFS file system.
- Eliminates a performance problem when a node, acting as CFS server of an NFS client file system, is write-appending to an external NFS server.
- Fixes a timing window during asynchronous reads on a CFS client.
- Fixes cfsmgr to properly return a failure status when a relocation request has failed.
- Fixes a race condition where stale name cache entries allow file access after file unlink.
- Fixes a panic that may occur during an unmount.
- Fixes an internal problem in the kernel's AdvFS, UFS, and NFS file systems where extended attributes with extremely long names, greater than 247 characters, could not be set on files. The new limit is 254 + a Null string terminator.
- Corrects a problem where a CFS lookup for a mount could leave stale state behind that could adversely affect subsequent NFS operations.

Number: Patch 87.00

Abstract: Fixes a panic that occurs on a booting node

State: Supersedes Patches 24.00, 43.00, 75.00, 77.00

- Addresses an assertion caused by a bad user pointer passed to the kernel via sys_call.
- Corrects a condition that causes a node to hang during testing the of Memory Channel cable pulls. A cluster member sometimes hangs when a Memory Channel cable is pulled, the node is taken down, the cable is plugged back in, and the node is rebooted.
- Increases performance by reducing the lock miss rate in the ics_mct_llnode_info_lock.
- Addresses a panic that occurs on a booting node.
- Addresses a panic that may occur when a node is joining the cluster. A node recognizing the joining node panics while it is trying to establish a preboot channel connection with the peer node, causing the folling message to be displayed on the console or in /var/adm/messages:

  panic (cpu x): ics_mct: rx conn 3

Number: Patch 89.00

Abstract: Fix for CAA core dumping problem

State: Supersedes Patch 33.00, 79.00

- Addresses an error "caa_register -u" produces with no balance data.
- Corrects a problem with resource inaccessibility if the hosting member crashes during a remote caa_stop operation.
- Fixes a problem in which CAA dumps core when trying to deal with cluster member ID 63.
- Fixes a problem in which CAAD might dump core due to a race condition when multiple events to which it subscribes arrive simultaneously.