

Tru64 UNIX and TruCluster Server Version 5.1A

Patch Summary and Release Notes for Patch Kit 4

February 2003

This manual describes the contents of Patch Kit 4 for the 5.1A 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.

© 2003 Hewlett-Packard Company

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

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

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

31-Jan-03

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	New Russian Keyboard — Patch 1197.00 (New)	1-1
1.1.3	Problem Seen on Systems with Smart Array Controller (New) ...	1-2
1.1.4	Problem with lockmode=4 on AlphaServers with QLogic SCSI Disk Controllers	1-2
1.1.5	Updates to sys_check	1-2
1.1.5.1	sys_check Version 125 Web Kit	1-2
1.1.6	Support for SDLT160 Tape Device	1-3
1.1.7	envconfig(8) Update — Tru64 UNIX Patch 156.00	1-4
1.1.8	Information for Tru64 UNIX Patches 1146.00 and 1148.00	1-4
1.1.9	Information for Tru64 UNIX Patch 252.00	1-4
1.1.10	Information for Tru64 UNIX Patch 1134.00	1-4
1.1.10.1	Updates to sh, csh, and ksh	1-5
1.1.10.2	sh noclobber Option and > , >> Constructs Added	1-5
1.1.10.3	ksh noclobber Behavior Clarified	1-5
1.1.10.4	csh noclobber Behavior Clarified	1-5
1.1.10.5	Updated mkdir System Call and Command	1-5
1.1.11	Reference Page Updates — Tru64 UNIX Patch 1107.00	1-6
1.1.12	Information for Tru64 UNIX Patch 1367.00	1-6
1.1.12.1	Enabling the /dev/poll Function	1-6
1.1.12.2	Removal of Version-Switched Patch	1-6
1.1.12.3	New ee Attribute	1-7
1.1.12.4	lag(7) and lagconfig(8) Reference Page Updates	1-7
1.1.12.5	wol(8) Reference Page Update	1-7
1.2	Summary of Base Operating System Patches	1-7

2 TruCluster Server Patches

2.1	Release Notes	2-1
2.1.1	Required Storage Space	2-1
2.1.2	Updates for Rolling Upgrade Procedures	2-1
2.1.2.1	Unrecoverable Failure Procedure	2-1
2.1.2.2	During Rolling Patch, Do Not Add or Delete OSF, TCR, IOS, or OSH Subsets	2-2
2.1.2.3	Undoing a Rolling Patch	2-2
2.1.2.4	Ignore Message About Missing ladebug.cat File During Rolling Upgrade	2-2
2.1.2.5	clu_upgrade undo of Install Stage Can Result in Incorrect File Permissions	2-2
2.1.2.6	Missing Entry Messages Can Be Ignored During Rolling Patch	2-3
2.1.2.7	Relocating AutoFS During a Rolling Upgrade on a Cluster ..	2-3

2.1.3	When Taking a Cluster Member to Single-User Mode, First Halt the Member	2-4
2.1.4	Additional Steps Required When Installing Patches Before Cluster Creation	2-4
2.1.5	Problems with clu_upgrade switch Stage	2-5
2.1.6	Cluster Information for Tru64 UNIX Patch 1367.00	2-5
2.1.7	Change to gated Restriction — TruCluster Patch 210.00	2-5
2.1.8	Information for TruCluster Patch 272.00	2-7
2.1.8.1	Enablers for EVM	2-7
2.1.8.2	Rolling Upgrade Version Switch	2-7
2.1.8.3	Restrictions Removed	2-7
2.1.9	CAA and Datastore — TruCluster Patch 242.00	2-8
2.2	Summary of TruCluster Software Patches	2-8

A Revised Reference Pages

A.1	envconfig(8) Update	A-1
A.2	sys_check(8) Update	A-3
A.3	sys_attrs_netrain(5), nifftmt(7), niffconfig(8), and ifconfig(8) Updates	A-9
A.3.1	sys_attrs_netrain(5)	A-9
A.3.2	nifftmt(7)	A-10
A.3.3	niffconfig(8)	A-11
A.3.4	ifconfig(8)	A-12
A.4	lag(7) Update	A-12
A.5	lagconfig(8) Update	A-15
A.6	wol(8) Update	A-18

About This Manual

This manual contains information specific to Patch Kit 4 of the Tru64 UNIX operating system and TruCluster Server software products for Version 5.1A. It provides a list of the patches contained in each kit and describes the information you need to know when installing specific 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.

Appendix A Provides changes to reference pages that occur as a result of 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.
- *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 the following 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*, the *Patch Summary and Release Notes* for a specific patch kit, and other patch-related documentation:

<http://www.tru64unix.compaq.com/docs/patch/>

- To visit our main support page:

<http://www.compaq.com/support/index.shtml>

- To visit the Tru64 UNIX homepage:

<http://www.tru64unix.compaq.com/>

Reader's Comments

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

You can send your comments in the following ways:

- Fax: 603-884-0120 Attn: UBPG Publications, ZK03-3/Y32
- Internet electronic mail:

readers_comment@zk3.dec.com

A Reader's Comment form is located on your system in the following location:
`/usr/doc/readers_comment.txt`

- Mail:

Hewlett-Packard Company
UBPG Publications Manager
ZK03-3/Y32
110 Spit Brook Road
Nashua, NH 03062-9987

Please include the following information along with your comments:

- The full title of this document.
- The section numbers and page numbers (if viewing in PDF format) 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.

Tru64 UNIX Patches

This chapter provides information about the patches included in Patch Kit 4 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 purpose of the Tru64 UNIX patches included in this kit.

Tru64 UNIX patch kits are cumulative. For this kit, this means that the patches and related documentation from patch kits 1 through 3 are included, along with patches that are new to this kit. To aid you in using this document, release notes that are new with this release are listed as (New) in the section head. The beginning of Section 1.2 provides a key for understanding the history of individual patches.

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 successfully install the base component of the patch kit:

- Approximately 71 MB of storage space in `/var/adm/patch/backup` may be required for archived original files if you choose to install and revert all patches. See the *Patch Kit Installation Instructions* for more information.
- Approximately 73 MB of storage space in `/var/adm/patch` may be required for original files if you choose to install and revert all patches. See the *Patch Kit Installation Instructions* for more information.
- Approximately 1346 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 2.1.1 for information on space needed for the TruCluster Server patches.

1.1.2 New Russian Keyboard — Patch 1197.00 (New)

The new Russian 3R-LKQ48–BT keyboard, for which Patch 1197.00 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,           0 ],
//     symbols[Group2]=3D [   Ukrainian_i,   Ukrainian_I ]
// };
key <AD09> {
    symbols[Group1]=3D [           o,           0 ],
    symbols[Group2]=3D [ Cyrillic_shcha, Cyrillic_SHCHA ]
};
```

1.1.3 Problem Seen on Systems with Smart Array Controller (New)

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
```

1.1.4 Problem with lockmode=4 on AlphaServers with QLogic SCSI Disk Controllers

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

1.1.5 Updates to `sys_check`

This section describes updates to the `sys_check` command. See Section A.2 for the revised `sys_check` reference page.

1.1.5.1 `sys_check` Version 125 Web Kit

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

This patch kit contains `sys_check` Version 124. If you have already installed the `sys_check` Version 125 web kit or higher, then installing this patch kit will downgrade the version of `sys_check` that is being used by the system. However,

you can easily set the system default back to the version of `sys_check` that you downloaded from the web by using the `/usr/sbin/use_sys_check` script. For example, type **use_sys_check 125** at the command line prompt to set `sys_check` Version 125 as the system default.

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

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

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

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

1.1.6 Support for SDLT160 Tape Device

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

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

```
scsi_density_table_size = 0x4a

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

SCSIDEVICE
#
# Matches SDLT320
#
Type = tape
Name = "COMPAQ" "SDLT320"
#
#
PARAMETERS:
TypeSubClass      = tk
TagQueueDepth     = 0
MaxTransferSize   = 0x0fffffb      # (16MB - 4)
ReadyTimeSeconds  = 120             # seconds

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

DENSITY:
#
DensityNumber = 1,5
```

```

DensityCode = default
CompressionCode = 0x1
Buffered = 0x1

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

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

```

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

1.1.7 envconfig(8) Update — Tru64 UNIX Patch 156.00

Patch 156.00 changes the `envconfig` command. See Section A.1 for an updated `envconfig(8)` reference page.

1.1.8 Information for Tru64 UNIX Patches 1146.00 and 1148.00

Patches 1146.00 and 1148.00 deliver version V2.0-094d of the `libots3` library. If your system has the Compaq FORTRAN Compiler, the Developer's Tool Kit (DTK) (OTABASE subset), or a patch that installs a newer version of this library, do not apply this patch. If a new revision of the `libots3` library is already installed on your system, and you install this patch, you will receive the following informational message:

```

Problem installing:

- Tru64_UNIX_V5.1A / Threads Patches

  Patch 00xxx.00 - Shared libots3 library fix

  ./usr/shlib/libots3.so:

    is installed by:

      OTABASE212
    and cannot be replaced by this patch.

```

This patch will not be installed.

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

```

# what /usr/shlib/libots3.so libots3.so:

libots3.a      V2.0-094 GEM 27 Feb 2001

```

1.1.9 Information for Tru64 UNIX Patch 252.00

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

1.1.10 Information for Tru64 UNIX Patch 1134.00

This section contains release notes for Patch 1134.00.

1.1.10.1 Updates to sh, csh, and ksh

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

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

```
Unable to create temporary file
```

1.1.10.2 sh noclobber Option and >| , >>| Constructs Added

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

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

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

1.1.10.3 ksh noclobber Behavior Clarified

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

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

1.1.10.4 csh noclobber Behavior Clarified

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

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

1.1.10.5 Updated mkdir System Call and Command

This kit reverts the `mkdir` system call, and thus the `mkdir` command, to its Tru64 UNIX Version 4.n behavior with respect to symbolic links. For the unusual case

where a symbolic link is used as the very last element of a `mkdir` path, the `mkdir` system call now returns an error rather than create the target.

If you want `mkdir` to follow the symbolic link you can do so by making the last character of the `mkdir` pathname a slash. For example, if `/var/tmp/foo` is a symbolic link to `/usr/xxx`, which does not exist, then `mkdir("/var/tmp/foo", 0755)` will return an error but `mkdir("var/tmp/foo/", 0755)` will create `/usr/xxx`.

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

1.1.11 Reference Page Updates — Tru64 UNIX Patch 1107.00

The installation of Patch 1107.00 results in changes to the `sys_attrs_netrain(5)`, `nifftmt(7)`, `niffconfig(8)`, and `ifconfig(8)` reference pages. To review these changes see the following sections:

- `sys_attrs_netrain(5)` See Section A.3.1.
- `nifftmt(7)` See Section A.3.2.
- `niffconfig(8)` See Section A.3.3.
- `ifconfig(8)` See Section A.3.4.

1.1.12 Information for Tru64 UNIX Patch 1367.00

This section contains release notes for Tru64 UNIX Patch 1367.00.

1.1.12.1 Enabling the `/dev/poll` Function

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

1. Change your directory to `/dev`:

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

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

1.1.12.2 Removal of Version-Switched Patch

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

Note

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

This script must be run by root in multiuser mode after completing the rolling upgrade that installed the patch and before starting another rolling upgrade. The final removal of the patch can only be accomplished by rebooting the system or cluster after this script completes its

processing. This script will offer to shut down your system or cluster at the end of its processing. If you choose to wait, it is your responsibility to execute the shutdown of the system or cluster.

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

1.1.12.3 New ee Attribute

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

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

`link_check_interval`

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

Default value: 200 (2 seconds)

Minimum value: 10 (0.1 seconds)

Maximum value: 1000 (10 seconds)

1.1.12.4 lag(7) and lagconfig(8) Reference Page Updates

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

See the *Network Administration: Connections* manual for information on configuring link aggregation groups. See `lag(7)` and `lagconfig(8)` for more information about link aggregation. See Section A.4 for the updated `lag(7)` reference page and Section A.5 for the `lagconfig(8)` reference page.

1.1.12.5 wol(8) Reference Page Update

The installation of this patch changes the information in the `wol(8)` reference page. See Section A.6 for the revised `wol(8)` reference page.

1.2 Summary of Base Operating System Patches

This section provides brief descriptions of the patches in Patch Kit 4 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 3)**
Indicates a patch that was new in the previous Version 5.1A patch kit.
- **Existing**
Indicates a patch that existed in earlier Version 5.1A patch kits.

- Supersedes Patches ...

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

Number: Patch 65.00

Abstract: Fix for Compaq C compiler and Compaq driver

State: Existing

This patch fixes the following problems in the Compaq C compiler and Compiler driver:

- A compiler problem that caused a run-time failure in specific code that involved floating point arguments and varargs.
 - A problem in the driver that failed to produce an object file for a command such as file.s -o file.o.
 - A problem in the driver that would not allow a command line that contained only the -l<arg> library and no source or object files.
 - A problem in the driver that failed to produce an object file when no output file was specified on the command line.
-

Number: Patch 88.00

Abstract: Fix for cluster hang during boot

State: Supersedes Patches 29.00

This patch fixes a situation in which the second node in a cluster hangs upon boot while setting current time and date with ntpdate.

Number: Patch 86.00

Abstract: Fix for Korn shell hang

State: Existing

This patch fixes a problem where the Korn shell (ksh) could hang if you pasted a large number of commands to it when it was running in a terminal emulator window (such as an xterm).

Number: Patch 97.00

Abstract: Fix for vi editor core dump problem

State: Existing

This patch fixes a problem where the vi editor core dumps when it finds invalid syntax during a substitute operation.

Number: Patch 108.00

Abstract: Fixes a potential race deadlock

State: Existing

This patch fixes a potential race deadlock between vclean/ufs_reclaim and quotaon/quotaoff when quota is enabled.

Number: Patch 117.00

Abstract: Fix for evmget command

State: Existing

This patch fixes a situation in which the evmget command and the event log nightly cleanup operation may fail with an "arg list too long" message.

Number: Patch 123.00

Abstract: Corrects a memory leak in the XTI socket code

State: Existing

This patch corrects a memory leak in the XTI socket code.

Number: Patch 136.00

Abstract: Fix for incorrect POSIX 4 message queues behavior

State: Existing

POSIX 4 message queue behavior was not following the standard and was returning unique message descriptors.

Number: Patch 138.00

Abstract: Static librt library fix for POSIX 4 message queues

State: Existing

POSIX 4 message queue behavior was not following the standard and returning unique message descriptors.

Number: Patch 141.00

Abstract: Security (SSRT0638U)

State: Existing (Kit 3)

This patch:

- Allows the dxsetacl utility to delete access ACLs.
 - Fixes a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of root directory compromise via lpr using X11.
-

Number: Patch 143.00

Abstract: Allows dxsetacl utility to delete access ACLs

State: Existing

This patch allows the dxsetacl utility to delete access ACLs.

Number: Patch 145.00

Abstract: Security (SSRT0638U)

State: Existing

This patch fixes a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of root directory compromise via lpr using X11.

Number: Patch 154.00

Abstract: Security (SSRT0682U)

State: Existing

This patch 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.

Number: Patch 156.00

Abstract: Fixes problems which prevented envmond from starting

State: Existing

This patch fixes problems which prevented envmond from starting.

Number: Patch 158.00

Abstract: Fix for Spike postlink optimizer

State: Existing

This patch fixes a problem where Spike may fail to delete the low instruction of a pair of related instructions, causing it to abort with a run-time error.

Number: Patch 169.00

Abstract: Fixes a problem in latsetup

State: Existing

This patch fixes a problem in latsetup when the directory /dev/lat is not found.

Number: Patch 171.00

Abstract: Fixes a problem in diskconfig

State: Existing

This fixes a problem in diskconfig where partitions with an offset and size of zero cannot be selected. It also fixes a problem where overlapping partitions cannot be adjusted if the existing partitions are not in alphabetical order.

Number: Patch 173.00

Abstract: Fix for ELSA Gloria Synergy, PS4D10, JIB graphic card

State: Existing

This patch fixes a problem where, on the ELSA Gloria Synergy, PS4D10, and JIB graphic cards, the cursor position is not being updated properly. The placement of the cursor is one request behind.

Number: Patch 185.00

Abstract: Corrects a problem in the rdist utility

State: Existing

This patch corrects a problem in the rdist utility which was causing segmentation faults on files with more than one link.

Number: Patch 187.00

Abstract: Fixes a volrecover error

State: Existing

This patch fixes a volrecover error of "Cannot refetch volume" when volumes exist only in a non-rootdg diskgroup.

Number: Patch 189.00

Abstract: Fix for no rerouting problem on a CFS server

State: Existing

This patch fixes a problem where pulling the network cable on one node acting as a CFS server in a cluster causes no rerouting to occur.

Number: Patch 195.00

Abstract: BPF default packet filter may cause system panic

State: Existing

This patch corrects a problem which could result in a system panic on close() if the BPF default packet filter is in use.

Number: Patch 210.00

Abstract: Fixes problems with X server X Image Extension (XIE)

State: Existing

This patch fixes problems with the X server X Image Extension (XIE).

Number: Patch 212.00

Abstract: Fixes a problem of the ATM setup script failing

State: Existing

This patch fixes a problem of the ATM setup script failing when configuring an elan if the lane subsystem is not loaded.

Number: Patch 220.00

Abstract: Corrects a problem with the NIFF daemon

State: Existing

This patch corrects a problem where the NIFF daemon (niffd) would exit if its connection to the EVM daemon (evmd) failed, as in the case of an EVM daemon restart.

Number: Patch 222.00

Abstract: Fix for mv command

State: Existing

This patch fixes a problem where the mv command will not perform a move if the inode of the file is the same as the inode of the destination directory, even though the file and directory are on different file systems.

Number: Patch 224.00

Abstract: The joint server may fail to clean up its lock files

State: Existing

The patch fixes a problem where joint may fail to clean up its lock files in /var/join.

Number: Patch 234.00

Abstract: Adds support for Persistent Reserve for HSV110

State: Existing

This patch updates the SCSI CAM Utility Program /sbin/scu by adding support for Persistent Reserve for HSV110 as well as the display of 128-bit WWIDS.

Number: Patch 238.00

Abstract: Fix for dxsetacl utility

State: Existing

This patch allows the dxsetacl utility to delete access ACLs.

Number: 248.00

Abstract: Fixes a problem in the strtod routine

State: Existing

This patch:

- Fixes a problem in which strtod() was returning different outputs for the same input.
 - Fixes a problem in which the tan() function was returning the wrong results.
-

Number: Patch 252.00

Abstract: Adds Essential Services Monitor daemon (esmd)

State: Existing (Kit 3)

This patch provides enablers for the Compaq Database Utility.

Number: Patch 259.00

Abstract: Removes extraneous header comments

State: Existing

This patch removes extraneous history edit comments from exported DECthreads header files.

Number: Patch 269.00

Abstract: Improves user control of clu_mibs

State: Existing

This patch move the control of the start and stop of the clu_mibs agent from /sbin/init.d/clu_max script to /sbin/init.d/snmpd script.

Number: Patch 281.00

Abstract: Fix for NHD kit installations

State: Existing

This patch corrects a problem that occurs during the installation of a New Hardware Device (NHD) kit, in which the version.id file was not properly referenced thereby causing the installation to fail.

Number: Patch 311.00

Abstract: Quick Setup erroneously reports daemons do not start

State: Existing

This patch fixes a problem that occurs on some systems, notably an AlphaStation DS10 system, in which Quick Setup may erroneously report that some daemons did not start and subsequent attempts generate other error messages appear that report duplicate host names.

Number: Patch 327.00

Abstract: Fixes C++ incompatibility

State: Existing

This patch fixes C++ incompatibility in three files in /usr/include/alpha/hal/ and one file in /usr/include/io/common/.

Number: Patch 414.00

Abstract: Fixes a problem in stdio.h

State: Existing

This patch fixes a problem in <stdio.h> where the interface renaming conditionals for fgetpos() and fsetpos() were mismatched. It also fixes a problem in <sys/timeb.h> where the ftime() prototype was not available in the default compilation name space.

Number: Patch 416.00

Abstract: Fixes a problem in sys/timeb.h

State: Existing

This patch fixes a problem in <stdio.h> where the interface renaming conditionals for fgetpos() and fsetpos() were mismatched. It also fixes a problem in <sys/timeb.h> where the ftime() prototype was not available in the default compilation name space.

Number: Patch 428.00

Abstract: Fix for evmwatch termination problem

State: Supersedes Patches 164.00 and 257.00

This patch:

- Resolves a memory leak and a filtering issue in the Event Manager, and allows the evmwatch utility to reconnect automatically if evmd fails and is restarted.
 - Fixes a problem in which binary error log (binlog) events posted by the EMX FibreChannel driver and the system console are reported incorrectly by the Event Manager, EVM.
 - Resolves an issue which can cause an Event Manager (EVM) client or the EVM daemon to core dump under rare circumstances.
-

Number: Patch 438.00

Abstract: Fix for dxproctuner utility

State: Existing

This patch fixes a problem in dxproctuner where the process information is not displayed when there is a double quotation mark followed by any other character in the command column.

Number: Patch 446.00

Abstract: Security (SSRT1-40U, SSRT1-41U, SSRT1-42U, SSRT1-45U)

State: New

This patch fixes 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 453.00

Abstract: Fix for dtgreet application

State: Existing

This patch fixes a problem in which after installing DCE, enabling SIA would cause a core dump and the greeter window does not come up.

Number: Patch 455.00

Abstract: Fix for lsmsa product

State: Existing

This patch addresses a problem in the display of disk controller to disk hierarchy by the lsmsa product.

Number: Patch 457.00

Abstract: Fix for broken symbolic links in /usr/lib/X11

State: Existing

This patch fixes a problem where three symbolic links in /usr/lib/X11 pointed to nonexistent directories.

Number: Patch 459.00

Abstract: Symbolic links point to nonexistent directories

State: Existing

This patch fixes a problem in which three symbolic links in /usr/lib/X11 point to nonexistent directories.

Number: Patch 465.00

Abstract: Fix for Elsa Gloria Comet card

State: Existing

This patch fixes a problem in which the Elsa Gloria Comet card does not correctly draw nested shaded boxes or anything similar.

Number: Patch 467.00

Abstract: Fix for accessx beeping functionality

State: Existing

This patch fixes a problem in which a beep does not occur when requested when the toggle keys option is enabled via accessx.

Number: Patch 469.00

Abstract: Fixes a problem in uucp

State: Existing

This patch fixes a problem in uucp. uucp between two Tru64 UNIX boxes hangs when a uucp failure occurs.

Number: Patch 481.00

Abstract: Security (SSRT1-40U, SSRT1-41U, SSRT1-42U, SSRT1-45U)

State: Existing

This patch fixes 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 485.00

Abstract: Fix for C++ compile problem

State: Existing

This patch fixes a C++ compile problem in /usr/include/X11/Xlib.h.

Number: Patch 490.00

Abstract: Fix for class scheduler failure

State: Existing (Kit 3)

This patch:

- Fixes a class scheduler semaphore race condition.
 - Causes the automatic detection a nonexistent semaphore and allocates a new one, thus allowing the class scheduler to proceed without interruption. The class scheduler depends on semaphores to protect its database from simultaneous updates.
-

Number: Patch 500.00

Abstract: Fixes libXm.so incompatibility

State: Existing

This patch fixes an libXm.so incompatibility.

Number: Patch 506.00

Abstract: Fix for rdump command

State: Existing

This patch causes the rdump command to properly dump data onto remote tape devices without receiving the SIGSEGV and dumping core.

Number: Patch 517.00

Abstract: Fix for LSM resynchronization problem

State: Existing

This patch corrects a problem with a mirrored LSM volume, with dirty region logging (DRL) enabled, still doing a full resynchronization during the first recovery after an unclean shutdown.

Number: Patch 519.00

Abstract: Fixes the C++ incompatibility with pwrmgr.h

State: Existing

This patch fixes the C++ incompatibility of /usr/include/dec/pwrmgr/pwrmgr.h.

Number: Patch 525.00

Abstract: Security (SSRT0779U)

State: Existing

This patch fixes a potential security vulnerability where, under certain circumstances, SNMP services can stop functioning.

Number: Patch 527.00

Abstract: Security (SSRT0779U)

State: Existing

This patch fixes a potential security vulnerability where, under certain circumstances, SNMP services can stop functioning.

Number: Patch 531.00

Abstract: Fix for KMF caused by malformed IPv4-in-IPv4 packets

State: Existing (Kit 3)

This patch:

- Corrects a condition in which a system configured with the IPTUNNEL kernel option will crash if it receives a corrupted IPv6-in-IPv4 packet, even if the system is not running IPv6. The system will panic with the message "kernel memory fault in ip6ip4_input()."
 - Fixes a kernel memory fault caused by malformed IPv4-in-IPv4 packets.
-

Number: Patch 533.00

Abstract: Fix for od command

State: Existing

This patch fixes a problem in which an invalid character sequence causes the od command to hang or display a partial character.

Number: Patch 535.00

Abstract: Fix for balance utility

State: Existing

This patch fixes problem in which the balance utility terminated before balancing the whole domain when the domain was very large (>4 GB).

Number: Patch 537.00

Abstract: Security (SSRT1-40U, SSRT1-41U, SSRT1-42U, SSRT1-48U)

State: Existing (Kit 3)

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

Number: Patch 545.00

Abstract: Fixes EVMs periodic channel monitoring function

State: Existing

This patch fixes a problem in which the Event Manager's channel monitoring function is temporarily disabled if the evmreload command is run.

Number: Patch 551.00

Abstract: Fixes an ATM signaling problem

State: Existing

This patch fixes an ATM signaling problem.

Number: Patch 553.00

Abstract: EVM daemon fails to find user-defined templates

State: Existing

This patch resolves a problem with the Event Manager (EVM) where user-defined events are not posted in a semirolled cluster. The Event Manager daemon fails to find user-defined templates in /usr/share/evm/templates/local on nonupgraded nodes in a semirolled cluster.

Number: Patch 565.00

Abstract: Enabler for Compaq Database Utility

State: Supersedes Patches 179.00, 292.00

This patch provides enablers for the Compaq Database Utility.

Number: Patch 569.00

Abstract: CD Mastering Software

State: Existing

This patch provides that CD Mastering Software for DS25 systems, which do not include a floppy drive, but have a CD-ROM burner instead. In order to write to this device, CD Mastering Software is required. It is made up of mkisofs and cdrecord software.

Number: Patch 571.00

Abstract: The savecore command prematurely terminates crash dump recovery

State: Existing

This patch corrects a problem where the savecore command may prematurely terminate crash dump recovery on partitions larger than 4 GB.

Number: Patch 578.00

Abstract: Fix for zdump utility

State: Existing

This patch:

- Fixes a problem in the zdump utility when time zone file names are specified as arguments without leading colons (:).
 - Fixes a regression in the -v output to display the current time.
-

Number: Patch 580.00

Abstract: Extended Visual Information returns incorrect info

State: Existing

This patch fixes a problem in which the X server's Extended Visual Information (EVI) extension was returning incorrect information.

Number: Patch 588.00

Abstract: Fix for NS record syntax in named.local file

State: Existing

This patch fixes the NS record syntax in a named.local file.

Number: Patch 691.00

Abstract: Fix for atexit and pthread_prefork handler crashes

State: Existing (Kit 3)

This patch fixes a problem with atexit() or pthread_atfork() handlers in shared libraries. An application will crash when handlers in shared libraries are called after the libraries are dlosed and unmapped.

Number: Patch 693.00

Abstract: Security (SSRT1-40U, SSRT1-41U, SSRT1-42U, SSRT1-45U)

State: Existing (Kit 3)

This patch fixes 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 699.00

Abstract: Fix for kdbx

State: Existing (Kit 3)

This patch fixes a premature termination of the ofile kdbx extension, warning messages in various kdbx extensions, and token length warnings when kdbx is invoked.

Number: Patch 701.00

Abstract: Security (SSRT1-40U, SSRT1-41U, SSRT1-42U, SSRT1-45U)

State: Existing (Kit 3)

This patch adds the mktemp(1) reference page for the mktemp command.

Number: Patch 705.00

Abstract: Security (SSRT1-40U, SSRT1-41U, SSRT1-42U, SSRT1-45U)

State: Existing (Kit 3)

This patch updates the mktemp(3) reference page with changed information regarding the mktemp() and mkstemp() routines, and adds information about the mkdtemp() and mkstemp() libc routines.

Number: Patch 711.00

Abstract: Fix for shfragbf

State: Existing

This patch clarifies the output of shfragbf, an AdvFS utility.

Number: Patch 713.00

Abstract: Fix for rcinet script

State: Existing

This patch prevents the system from hanging when the rcinet script is used by correcting the order in which NetRAIN-related services are started and stopped.

Number: Patch 718.00

Abstract: Fix for traceroute command

State: Existing (Kit 3)

This patch corrects a problem where traceroute sometimes failed to provide responses and finish a trace when the destination host name was given on the command line.

Number: Patch 720.00

Patch: Fix for assembler problems

State: Existing (Kit 3)

This patch, shipped as Version 3.06.08 of the Tru64 UNIX Assembler, resolves three assembler problems related to the following:

- The generation of an incorrect symbol table which can cause om to fail.
 - The improper reordering of an instruction which restores the stack pointer when assembling with optimization active.
 - The generation of a .ident string without a terminating NULL.
 - The generation of an invalid optimization when a load instruction specifies a target register and a base register that are the same.
-

Number: Patch 731.00

Abstract: (SSRT1-40U, SSRT1-41U, SSRT1-42U, SSRT1-45U)

State: Existing (Kit 3)

This patch adds the safe_open(3) reference page for the safe_open() routine in libc.

Number: Patch 733.00

Abstract: Fix for Memory Channel driver problem

State: Existing (Kit 3)

This patch addresses a situation in which the Memory Channel device shuts down if too many state change interrupts are received.

Number: Patch 739.00

Abstract: Fixes a problem with os_mibs

State: Existing (Kit 3)

This patch fixes a problem with os_mibs that could cause the application to consume an excessive amount of CPU time.

Number: Patch 741.00

Abstract: Enhancement to savemeta script

State: Existing (Kit 3)

This patch enhances the capability of savemeta script as follows:

- Allows the script to be used in single user mode on a corrupt /usr domain.
 - Causes all errors to return 1.
-

Number: Patch 744.00

Abstract: Shared library fix for libaio

State: Existing (Kit 3)

This patch:

- Fixes a rarely seen memory fault in libaio during aio_cancel().
 - Adds support for NEW_OPEN_MAX_SYSTEM (64K) file descriptors to libaio.
 - Prevents thread blocking forever when both libaio and libaio_raw are linked into the same image.
 - Closes an aio_read()/aio_cancel() race condition.
-

Number: Patch 747.00

Abstract: Static library fix for libaio

State: Existing (Kit 3)

This patch:

- Fixes a rarely seen memory fault in libaio during aio_cancel().
 - Adds support for NEW_OPEN_MAX_SYSTEM (64K) file descriptors to libaio.
 - Prevents thread blocking forever when both libaio and libaio_raw are linked into the same image.
 - Closes an aio_read()/aio_cancel() race condition.
-

Number: Patch 749.00

Abstract: Modification to secconfig suitlet

State: Existing (Kit 3)

This patch makes the customize database option available when using secconfig for shadow passwords.

Number: Patch 751.00

Abstract: Security (SSRT0818U)

State: Existing (Kit 3)

This patch 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.

Number: Patch 753.00

Abstract: Security (SSRT0818U)

State: Existing (Kit 3)

This patch 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.

Number: Patch 755.00

Abstract: Security (SSRT0818U)

State: Existing (Kit 3)

This patch 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.

Number: Patch 757.00

Abstract: Fix for script command

State: Existing (Kit 3)

This patch corrects a problem in which script would hang upon exit in a dfs configuration.

Number: Patch 759.00

Abstract: System Mgmt Station detects failing PCI adapters

State: Existing (Kit 3)

This patch provides the ability for the System Management Station to render PCI adapters with a warning or failed representation when they are in the indicted state. This is in addition to the previous ability to render CPUs that are in the indicted state.

Number: Patch 761.00

Abstract: Fixes a problem in the mwm window manager

State: Existing (Kit 3)

This patch fixes a problem in the mwm window manager in which double-click actions are performed on the second button press instead of the second button release, thus causing the second button release event to be sent to any underlying window.

Number: Patch 763.00

Abstract: Security (SSRT1-40U, SSRT1-41U, SSRT1-42U, SSRT1-45U)

State: Existing (Kit 3)

This patch adds the dirclean(8) reference page for the /usr/sbin/dirclean utility.

Number: Patch 765.00

Abstract: Provides the poll reference page

State: Existing (Kit 3)

This patch adds the poll(7) reference page for the /dev/poll driver.

Number: Patch 767.00

Abstract: Enhancement to fuser utility

State: Existing

This patch allows the fuser utility to display the reference option. This option indicates the type of reference made; for example: open, closed, unlinked, or mmapped.

Number: Patch 769.00

Abstract: Fix for su command

State: Existing (Kit 3)

This patch corrects the behavior of the su command so that the LOGNAME environment variable is changed to the target user when executed with the - option.

Number: Patch 771.00

Abstract: lsmsa incorrectly processing passwords

State: Existing (Kit 3)

This patch fixes a problem where lsmsa incorrectly processes passwords that are longer than eight characters. Anyone who tries to start the LSM GUI using a password of eight or more characters will be denied access.

Number: Patch 773.00

Abstract: Security (SSRT0795U)

State: Existing (Kit 3)

This patch fixes a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form where inetd may block incoming connections when scanned by nmap or other port scanners.

Number: Patch 779.00

Abstract: Fixes an xfs problem

State: Existing (Kit 3)

This patch fixes an xfs problem which causes a "QueryGlyphs failed" error in showfont.

Number: Patch 787.00

Abstract: Enhancement for siacfg utility

State: Existing (Kit 3)

This patch allows nonlocal SIA mechanisms, LDAP for example, to place their mechanism last in the list of mechanisms.

Number: Patch 789.00

Abstract: Fix for ppdof print filter core dump problem

State: Existing (Kit 3)

This patch corrects a problem where the filter can core dump when the banner jobname contains 132 characters.

Number: Patch 791.00

Abstract: Fix for salvage utility core dump problem

State: Existing (Kit 3)

This patch fixes a problem with the /sbin/advfs/salvage utility that could cause the utility to core dump.

Number: Patch 793.00

Abstract: Fix for startslip program problem

State: Existing (Kit 3)

This patch fixes a problem where startslip was not able to extract all the information from the acucap file.

Number: Patch 795.00

Abstract: A timing window can cause a hang in run_usr_cmd

State: Existing (Kit 3)

This patch fixes a problem in which a timing window can cause a hang in run_usr_cmd.

Number: Patch 801.00

Abstract: Fix for convuser utility

State: Existing (Kit 3)

This patch fixes a problem where, if a user was working in enhanced security and then switched to base security, the group and other read privileges would get stripped from /etc/passwd.

Number: Patch 803.00

Abstract: Fixes a problem in access_lines for libmld

State: Existing (Kit 3)

This patch fixes a problem that may cause the third command and other Atom-based instrumentation tools to fail.

Number: Patch 809.00

Abstract: Enables correctable error reporting from DTAGII

State: Existing (Kit 3)

This patch enables correctable error reporting from DTAGII chips on GS320/160/80 1.224Ghz CPU systems.

Number: Patch 817.00

Abstract: Fixes a simple_lock panic when using ATM

State: Existing (Kit 3)

This patch:

- Fixes a kernel memory fault when using ATM.
- Corrects a problem which could result in ATM/lane connection requests being dropped.
- Fixes a kernel memory fault when using ATM.
- Fixes a "simple_lock: time limit exceeded" panic when using ATM.

Number: Patch 824.00

Abstract: Provides the ckfsec reference page

State: Existing (Kit 3)

This patch delivers the ckfsec(1) reference page.

Number: Patch 826.00

Abstract: Security (SSRT0794U)

State: Existing (Kit 3)

This patch fixes a potential security vulnerability where, under certain circumstances, system integrity may be compromised.

Number: Patch 828.00

Abstract: Security (wc.symlink.002.spautils)

State: Existing (Kit 3)

This patch fixes a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of certain files in world-writable directories. This patch provides the ckfsec utility which can help detect such files.

Number: Patch 830.00

Abstract: Incompatibility between Java 1.1.x and Java 2 1.2.x

State: Existing (Kit 3)

This patch allows users who are using SysMan tools on the iPAQ, or who have Java 2 as their default Java version, to communicate with Tru64 UNIX V5.1A systems.

Number: Patch 832.00

Abstract: Update to exportfs reference page

State: Existing (Kit 3)

This patch updates the exportfs(2) reference page with changed information regarding the exportfsdata structure as a result of increasing a number of file systems that can be NFS mounted from 256 to 1024.

Number: Patch 840.00

Abstract: Definition of XtPending was changed

State: Existing (Kit 3)

This patch fixes a problem where the definition of the X Toolkit function XtPending() was changed in Tru64 UNIX V5.1A which caused some applications built on earlier versions of Tru64 UNIX to fail.

Number: Patch 848.00

Abstract: Security

State: New

This patch fixes 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 850.00

Abstract: Security

State: New

This patch fixes 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 852.00

Abstract: Security (SSRT1-40U, SSRT1-41U, SSRT1-42U, SSRT1-45U)

State: Supersedes Patch 92.00

This patch:

- Fixes a potential security vulnerability where, under certain circumstances, users can clobber temporary files created by shell commands and utilities, for example under /sbin, /usr/sbin, /usr/bin, and /etc.
- 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.

Number: Patch 992.00

Abstract: Fix for library call used to calculate disk size

State: Supersedes Patch 844.00

This patch:

- Fixes a problem that prevented AdvFS from working correctly with LSM volumes between 1Tb and 2Tb.
- Causes mkfdmn and addvol to issue a warning if an attempt is made to use an LSM volume greater than 2Tb.
- Prevents addvol from adding invalid disks into a domain.

Number: Patch 994.00

Abstract: Fix for problems in dsfmgr

State: New

This patch fixes many small problems in dsfmgr.

Number: Patch 999.00

Abstract: Security (SSRT1-80U)

State: Supersedes Patches 71.00 and 997.00

This patch:

- Fixes several potential security vulnerabilities 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.
- 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.

Number: Patch 1002.00

Abstract: Security (SSRT1-80U)

State: Supersedes Patches 73.00 and 1000.00

This patch:

- Fixes a potential security vulnerability where, under certain circumstances, users can clobber temporary files created by shell commands and utilities (that is, under /sbin, /usr/sbin, /usr/bin, and /etc).
 - 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.
-

Number: Patch 1008.00

Abstract: Security (SSRT0767U, SSRT2251, SSRT2274, SSRT0788U)

State: Supersedes Patches 539.00, 208.00, 429.00, 430.00, 432.00, 695.00, 1003.00, 1004.00, 1005.00, 1006.00, 1283.00

This patch: :

- Fixes a potential security vulnerability where, under certain circumstances, system integrity may be compromised. The tttdserverd contains a potential buffer overflow that may allow unauthorized access.
- Fixes 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 a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of large values of command line arguments.
- Fixes a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of large values of ENVIRONMENT variables and command line arguments.
- Fixes a potential security vulnerability in CDE Subprocess Control Service (dtspcd), where 'dtspcd has a potential buffer overflow condition which may lead to unauthorized access.
- Fixes a problem where palette files are not from /etc/dt/palettes.
- Fixes a problem with dtprintinfo memory fault with long LANG value.
- Fixes a problem where a CDE session hangs at startup using localized .dt files located in ~/.dt/types directory.

Number: Patch 1011.00

Abstract: Security (SSRT0753U, SSRT0752U, SSRT0788U)

State: Supersedes Patches 433.00, 434.00, 436.00, 697.00, 1009.00

This patch: :

-
- Fixes several potential security vulnerabilities where, under certain circumstances, system integrity may be compromised. These may be in the form of large values of command line arguments.
- Fixes a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of large values of ENVIRONMENT variables and command line arguments.
- Fixes a potential security vulnerability in CDE Subprocess Control Service (dtspcd), where 'dtspcd has a potential buffer overflow condition which may lead to unauthorized access.
- Fixes several potential security vulnerabilities where, under certain circumstances, system integrity may be compromised. These may be in the form of improper file access.
- Corrects a dtprintinfo memory fault problem with long LANG value.

Number: Patch 1020.00

Abstract: Security (SSRT2275)

State: Supersedes Patch 1018.00

This patch:

- Fixes 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 several potential security vulnerabilities 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.
-

Number: Patch 1023.00

Abstract: Fixes memory leak in Panoramix/Xinerama Extension

State: Supersedes Patches 98.00, 99.00, 101.00, 439.00, 441.00, 1021.00

This patch:

- Provides New Hardware Delivery V4 (NHD4) enables for future hardware support of a graphics device.
- Fixes an Xserver crash when using the GTK on systems using the Oxygen VX1 graphics card.
- Fixes an Xserver problem where, when Panoramix is enabled and using CDE, icons from dtfile cannot be seen on other than the left screen while being moved.
- Fixes a memory leak in the Panoramix/Xinerama Extension that could cause a process core dump.
- Fixes a problem with a Compaq Professional Workstation XP1000 667 MHz system with a PowerStorm 4D20 (PBXGB-CA) graphics card where fonts were sometimes drawn incorrectly.
- Fixes a problem where the X Window System XGetImage() function returns erroneous data for displays with a depth greater than 8 when running the Panoramix extension.
- Corrects XCopyPlane on the Oxygen VX1 graphics card to copy only the requested bitplane rather than all bitplanes.

Number: Patch 1025.00

Abstract: Security (SSRT2301)

State: New

This patch fixes a potential security vulnerability where, under certain circumstances, system integrity may be compromised through improper file access (overwriting of files). This potential vulnerability is in the form of a local security domain risk. The following potential security vulnerability has been corrected:

SSRT2301 uudecode (Severity - Medium)

Number: Patch 1027.00

Abstract: Security

State: New

This patch 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.

Number: Patch 1040.00

State: Supersedes Patches 3.00, 5.00, 113.00, 115.00, 119.00, 447.00, 479.00, 451.00, 449.00, 706.00, 707.00, 709.01, 1028.00, 1029.00, 1030.00, 1031.00, 1032.00, 1033.00, 1034.00, 1035.00, 1036.00, 1037.00, 1038.00

This patch:

- Fixes AutoFS problems as follows:
 - Eliminates error messages concerning property lists seen via certain utilities such as vdump.
 - Causes AutoFS auto-mounts to occur when utilities name intercept points defined via indirect map entries.
 - Fixes a deadlock that will occur in non-cluster systems when direct map entries are served locally.
- Fixes vdump command problems as follows:
 - Fixes a problem in which the command fails to flag compressed extended attributes records that are split across a vdump BLOCK boundary
 - Corrects a "Rewinding" message to avoid a segfault with Internationalized messages.

Patch 1040.00 Continued

- Fixes vrestore command problems as follows:
 - Fixes a problem in which the command fails to flag compressed extended attributes records that are split across a vdump BLOCK boundary.
 - Fixes a problem in which the command fails to set extended attributes due to confusion over selective restore of the file or directory associated. Also results in display of the error message "error setting extended attributes."
 - Fixes a problem in which the selective restore of hardlinked files is incomplete when they exist in different directories (fails to create directory for second occurrence of file with same inode number).
- Eliminates inefficient behavior by autofs when the top level directory of a direct hierarchical auto-mount map entry cannot be successfully mounted.
- Ensures that AutoFS correctly uses the mount options specified in auto-mount map entries with replicated servers.
- Fixes a problem where the tar -F (Fasttar) command ignores files named err but does not ignore files named errs and directories named SCCS and RCS.
- Corrects pax/tar/cpio to properly extract explicitly specified files.
- Corrects the behavior of several commands when used in conjunction with file systems that are locally served via AutoFS.
- Provides support for wildcards in Linux /etc/exports entries. Both AutoFS and Automount have been so enhanced.
- Fixes a problem where autofs and autofs mount daemons do not properly parse "*" in map files.
- Fixes a problem that prevents access to AutoFS file systems if ACLs are enabled.
- Allows auto-unmounts to succeed for direct map entries in which the key contains a symbolic link.
- Fixes a one byte gap/hole in the maximum size in the tar command before an extended header record is used (8589934591 (octal 7777777777)).
- Eliminates AutoFS auto-mount failures due to the receipt of unexpected signals while sleeping in the kernel.
- Corrects find -ls, which displays an incorrect number of blocks.
- Supports a cluster patch that limits DLM lock contention by AutoFS in a cluster.
- Fixes a defect that allows the possibility that certain AutoFS auto-mounted filesystems may be mounted more than once.
- Ensures that autofs mount will process all NIS-provided map files included via the "+" syntax. Currently, only the first one, at a given level of includes, will be correctly processed.
- Eliminates extraneous communication between an AutoFS client and an NFS server's mount daemon in some situations when replicated servers are specified in an auto-mount map entry.
- Allows AutoFS to correctly handle NIS-based auto-mount maps using the wildcard (*) key.
- Corrects two debugging/diagnostic messages in the AutoFS utilities
- Fixes a cluster-specific problem with AutoFS in which a race condition leads to a slowdown, which may lead to auto-mount failures.

Number: Patch 1042.00**Abstract:** Fix for libos shared library**State:** New

This patch fixes a problem where the X server's command line option to turn off VESA Display Power Management Signalling (-dpms) does not work.

Number: Patch 1044.00

Abstract: Fixes failures under high DMA resource utilization

State: New

This patch corrects a problem that can occur under high DMA resource utilization. If a request for DMA resources (via `dma_map_load()`) fails, a stale pointer to freed memory is returned to the requestor, setting up the potential for a double free of malloc'd memory or dereferencing through that pointer, resulting in a panic.

Number: Patch 1046.00

Abstract: Security

State: New

This patch fixes 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 1059.00

Abstract: Fix for Sysman Account Management tool

State: New (Supersedes Patch 1057.00)

This patch:

- Deals with errors that can be generated when (non-root) users, who lack a valid Home Directory, run the Sysman Account Management tool.
 - Deals with errors that are generated when there are +users or -users present in the `/etc/passwd` file and the Sysman Account Management Suitlet GUI tool is used to display or modify that +user or -user.
-

Number: Patch 1061.00

Abstract: Security (SSRT2208)

State: New

This patch fixes a potential security vulnerability which may allow non-privileged users to gain unauthorized (root) access. This may be in the form of local and remote security domain risks. This potential security vulnerability in routed has been corrected.

Number: Patch 1063.00

Abstract: Fixes hwmgr command to show path state correctly

State: New

This patch fixes the hwmgr command to show path state correctly.

Number: Patch 1065.00

Abstract: Fix for cp command

State: New (Supersedes Patch 160.00)

This patch fixes a problem in which the `cp` and `cat` commands produce different file sizes when reading from a tape device. It also corrects the `cp` command performance related to the problem.

Number: Patch 1073.00

Abstract: Security (SSRT2229)

State: New

This patch 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.

Number: Patch 1075.00

Abstract: Security

State: New

This patch fixes 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 1087.00

Abstract: Security

State: New

This patch fixes 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 1089.00

Abstract: Security

State: New

This patch fixes 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 1091.00

Abstract: Security (SSRT1-40U, SSRT1-41U, SSRT1-42U, SSRT1-45U)

State: New

This patch fixes 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 1093.00

Abstract: Security (SSRT2339, SSRT2339)

State: New

This patch fixes 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 1095.00

Abstract: Security

State: New

This patch fixes 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 1097.00

Abstract: Fix prevents "simple lock owned" panics

State: New

This patch prevents "simple lock owned" panics.

Number: Patch 1099.00

Abstract: Security

State: New

This patch fixes 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 1101.00

Abstract: Security

State: New

This patch fixes 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 1103.00

Abstract: Fix for verify command

State: New

This patch fixes a problem whereby the verify utility would core dump if it encountered a specific type of metadata inconsistency.

Number: Patch 1105.00

Abstract: Security (SSRT1-40U, SSRT1-41U, SSRT1-42U, SSRT1-45U)

State: Supersedes Patch 502.00

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

Number: Patch 1107.00

Abstract: Adds support to ifconfig application for IPv6 command

State: Supersedes Patches 232.00 and 504.00)

This patch:

- Fixes a problem in NetRAIN by causing NetRAIN interface creation to fail if any of the requested standby interfaces do not exist.
 - Provides for faster failover time, mainly by permitting timer values of less than 1 second to be configured.
 - Adds support to ifconfig application for the IPv6 command line argument ip6reachabletime.
-

Number: Patch 1109.00

Abstract: Security

State: New

This patch 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.

Number: Patch 1111.00

Abstract: Security

State: New

This patch fixes 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 1113.00

Abstract: Security

State: New

This patch 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.

Number: Patch 1115.00

Abstract: Fixes interop problem between curses.h and esnmp.h

State: New

This patch fixes an interoperability problem between the curses.h and esnmp.h header files.

Number: Patch 1117.00

Abstract: Security

State: New

This patch 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.

Number: Patch 1121.00

State: New

This patch allows non-Compaq C compilers to avoid name space pollution when getting information about a file.

Number: Patch 1123.00

Abstract: Security

State: New

This patch 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.

Number: Patch 1125.00

State: New

This patch fixes the display for the `hwmgr -show name` command to align properly for the name field.

Number: Patch 1127.00

Abstract: Security

State: New

This patch fixes 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 1129.00

Abstract: Update to `hwmgr` utility

State: New

This patch fixes a problem where, when using `hwmgr` to delete a component, the message "DELETE_COMMIT: Cannot fetch name" may be displayed on the console.

Number: Patch 1134.00

Abstract: Security (SSRT1-40U, SSRT1-41U, SSRT1-42U, SSRT1-45U)

State: Supersedes Patches 125.00, 309.00, 460.00, 461.00, 463.00, 734.00, 735.00, 737.00, 1130.00, 1131.00, 1132.00

This patch:

- Fixes several potential security vulnerabilities where, under certain circumstances, system integrity may be compromised when a buffer overflow occurs in the `ksh` or `sh` utilities. 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.
 - Fixes several potential security vulnerabilities where, under certain circumstances, system integrity may be compromised. These may be in the form of improper file access.
 - Fixes a problem in which `/usr/bin/ksh` hangs for certain scripts that contain the `wait` command.
 - Makes shell inline input files more secure.
 - Adds `sh noclobber` and new constructs.
 - Updates the `mkdir` system call.
 - Corrects a problem in which `ksh` fails to substitute the tilde (`~`) character for a user's home directory after making an assignment using the `"#"` or `"%"` characters.
 - Fixes a problem with `ksh` in when a `ksh` menu is started from within user's `.profile` file, `ksh` will not stop when the Telnet session is stopped.
 - Fixes an Asian language processing problem under the Korn shell.
 - Corrects a problem in which `sh` was using a high amount of CPU time.
 - Corrects a problem in which `sh` will not receive a `SIGSEGV` signal when the user runs the `type` utility with a file path greater than 69 characters.
-

Number: Patch 1136.00

Abstract: Modifies enablers for Enterprise Volume Manager

State: Supersedes Patches 197.00, 263.00, 313.00, 513.00

This patch modifies enablers for Enterprise Volume Managert.

Number: Patch 1138.00

Abstract: Modifies enablers for Enterprise Volume Manager

State: Supersedes Patches 201.00, 265.00, 317.00, 515.00

This patch modifies enablers for Enterprise Volume Manager.

Number: Patch 1140.00

Abstract: Modifies enablers for Enterprise Volume Manager

State: Supersedes Patches 199.00, 267.00, 315.00

This patch modifies enablers for Enterprise Volume Manager.

Number: Patch 1142.00

Abstract: Security (SSRT0792U)

State: New

This patch 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.

Number: Patch 1144.00

Abstract: Security

State: New

This patch fixes 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 1146.00

Abstract: Installs version V2.1-120 of libots3

State: Supersedes Patch 226.00

This patch installs version V2.1-120 of the /usr/lib/libots3.a and /usr/shlib/libots3.so libraries, which do the following:

- Fixes a problem in which the max threads clause for the SGI parallel interfaces is being ignored.
 - Fixes a problem in which an OpenMP thread may hang when reaching a critical region and all other threads are awaiting CVs.
 - Fixes a problem in which long-running OpenMP applications might overflow an internal libots3 counter, resulting in a breakdown of thread synchronization.
-

Number: Patch 1148.00

Abstract: Installs version V2.1-120 of libots3

State: Supersedes Patch 228.00

This patch installs version V2.1-120 of the /usr/lib/libots3.a and /usr/shlib/libots3.so libraries, which do the following:

- Fixes a problem in which the max threads clause for the SGI parallel interfaces is being ignored.
 - Fixes a problem in which an OpenMP thread may hang when reaching a critical region and all other threads are awaiting CVs.
 - Fixes a problem in which long-running OpenMP applications might overflow an internal libots3 counter, resulting in a breakdown of thread synchronization.
-

Number: Patch 1152.00

Abstract: Security

State: Supersedes Patches 584.00, 775.00

This patch:

- Fixes a memory leak problem in the Window Manager.
 - Fixes a problem in the dtwm window manager where double-click actions are performed on the second button press instead of the second button release, which causes the second button release event to be sent to any underlying window.
 - Fixes 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 1154.00

Abstract: Security

State: New

This patch fixes a potential security vulnerability where, under certain circumstances, system integrity may be compromised when a buffer overflow occurs in the binmail (also called mail) 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 1156.00

Abstract: Security

State: New

This patch 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.

Number: Patch 1158.00

Abstract: Security

State: Supersedes Patch 529.00

This patch:

- Fixes a problem where no shell message is displayed when trying to invoke the su command to a user other than root.
 - Fixes a potential security vulnerability where, under certain circumstances, system integrity may be compromised.
-

Number: Patch 1166.00

Abstract: Fixes consvar -s bootdef_dev failure with KZPCC

State: New

This patch:

- Fixes a problem where I/O greater than 4MB fails to KZPCC devices with the error ENODEV.
 - Fixes the problem of failed open calls to KZPCCs under heavy I/O.
 - Fixes consvar -s bootdef_dev failure with KZPCC.
-

Number: Patch 1168.00

Abstract: Fix for libdix shared library

State: Supersedes Patches 76.00 and 78.00

This patch:

- Fixes a problem that will cause the X server to hang on rare occasions. Except for the mouse, everything on the desktop appears frozen. Output from the ps command will show the X server using greater than 99% of the CPU time.
 - Fixes a problem that can cause CDE pop-up menus to appear on the wrong screen when running a multihead system with the Panoramix extension enabled.
 - Fixes a problem that causes the reversal of black and white colors on screens other than screen 0 of a multihead system.
-

Number: Patch 1170.00

Abstract: Security

State: Supersedes Patch 84.00

This patch:

- Fixes a problem in which shutting down the network would also shut down the cluster interconnect interface in a LAN cluster.
 - Fixes a potential security vulnerability which, under certain circumstances, could compromise system integrity. This may be in the form of improper file or privilege management.
-

Number: Patch 1172.00

Abstract: Fix for sysconfig utility

State: New

This patch fixes a problem in which the lines in the output stream from "sysconfig -Q" can be truncated.

Number: Patch 1175.00

Abstract: Security

State: New (Supersedes Patch 1173.00)

This patch:

- Fixes core dump problem when using the SysMan SUItlet to configure NFS daemons in curses mode.
- Fixes a potential security vulnerability which, under certain circumstances, could compromise system integrity. This may be in the form of improper file access.

Number: Patch 1177.00

Abstract: Security (SSRT0743U, SSRT2256)

State: Supersedes Patch 69.00

This patch:

- Fixes a potential security vulnerability which, under certain circumstances, could compromise system integrity. This may be in the form of improper file or privilege management.
- Fixes a potential security vulnerability where, under certain circumstances, system integrity may be compromised when a buffer overflow occurs in the ps' 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 1179.00

Abstract: Security

State: Supersedes Patch 106.00

This patch:

- Corrects the behavior of the sort(1) command which now checks for duplicates with -c -u and -k flags.
 - Fixes a potential security vulnerability which, under certain circumstances, could compromise system integrity. This may be in the form of improper file or privilege management.
-

Number: Patch 1181.00

Abstract: Security (SSRT1-40U, SSRT1-41U, SSRT1-42U, SSRT1-45U)

State: Supersedes Patches 109.00, 110.00, 112.00, 282.00, 284.00, 442.00, 444.00, 714.00, 716.00, 1047.00, 1048.00, 1049.00, 1050.00, 1051.00, 1052.00, 1053.00, 1054.00, 1056.00

This patch:

- Fixes several potential security vulnerabilities which, under certain circumstances, could compromise system integrity. These may be in the form of improper file access.
- Fixes a problem with the SysMan Station that causes incorrect state information to be displayed after a CPU has been indicted.
- Fixes possible deadlock conditions in the SysMan station daemon that might occur at daemon startup or during failover.
- Provides enablers for Packaged Database Solution.
- Corrects a problem in which objects in the Physical File system view do not have correct or updated properties.
- Corrects a problem in which the SysMan Station cannot launch commands on objects where an object attribute is part of the command.
- Fixes a problem where reconfiguration of network interface cards using SysMan makes the old IP address an IP alias. The new IP address now replaces the old IP address.
- Fixes the message catalogs in some legacy applications so that the proper pop-up error message is displayed at the appropriate time.
- Fixes a problem which could cause communications over a cluster interconnect to break when gigabit Ethernet cards are used as cluster interconnects. The problem could occur because the cards are shown in the netconfig suitlet as regular interface cards and a user may reconfigure this, potentially calling a break in the interconnect.
- Corrects a problem in which Quick Setup does not check for failure when safely creating a temporary file. Currently, it checks for failure and aborts the procedure if the action failed.

Number: Patch 1183.00

Abstract: Security

State: New

This patch fixes a potential security vulnerability which, under certain circumstances, could compromise system integrity. This may be in the form of improper file access.

Number: Patch 1185.00

Abstract: Security

State: New

This patch fixes a potential security vulnerability which, under certain circumstances, could compromise system integrity. This may be in the form of improper file or privilege management.

Number: Patch 1187.00

Abstract: Security

State: New

This patch fixes a potential security vulnerability which, under certain circumstances, could compromise system integrity. This may be in the form of improper file or privilege management.

Number: Patch 1190.00

Abstract: Fix for ld linker

State: Supersedes Patches 150.00, 555.00, 1188.00

This patch:

- Fixes a problem that causes the linker (ld) to crash when certain data alignment directives are used in the link.
- Fixes a problem with the datatype of the linker-defined `_fpdata` symbol.
- Fixes a problem in which the linker may corrupt the shared object registry file when `-update_registry` is specified with concurrent links.
- Fixes a linker error that occurs when the command `ld -update_registry /dev/null` is specified.
- Fixes a linker problem that may cause executables to fail with a segmentation violation when the address of an uninitialized data symbol in a shared library is used as the initial value of a global or static pointer variable.

Number: Patch 1192.00

Abstract: Security

State: New

This patch fixes a potential security vulnerability which, under certain circumstances, could compromise system integrity. This may be in the form of improper file access.

Number: Patch 1197.00

Abstract: Provides updated keyboard map for Russian keyboard

State: New

This patch provides an updated keyboard map for the Russian 3R-LKQ48-BT keyboard model.

Number: Patch 1200.00

Abstract: Performance tool failures on Sierra Clusters PFS

State: New (Supersedes Patch 1198.00)

This patch:

- Fixes a problems in which `prof -pixie -testcoverage` command sometimes reports invalid source line number ranges.
- Fixes performance tool failures on Sierra Clusters (PFS) Parallel File Systems.

Number: Patch 1202.00

Abstract: Security

State: New

This patch fixes a potential security vulnerability which, under certain circumstances, could compromise system integrity. This may be in the form of improper file access.

Number: Patch 1206.00

Abstract: Security

State: New (Supersedes Patches 1203.00 and 1204.00)

This patch fixes several potential security vulnerabilities which, under certain circumstances, could compromise system integrity. These may be in the form of improper file or privilege management.

Number: Patch 1208.00

Abstract: Security (SSRT1-40U, SSRT1-41U, SSRT1-42U, SSRT1-45U)

State: Supersedes Patch 573.00

This patch fixes several potential security vulnerabilities which, under certain circumstances, could compromise system integrity. These may be in the form of improper file access.

Number: Patch 1210.00

Abstract: Fixes a simple lock fault in the floppy driver

State: New

This patch fixes a simple lock panic in the floppy driver.

Number: Patch 1212.00

Abstract: Fixes and improves the mcutil program

State: New

This patch fixes and improves the mcutil program by correcting how bus resets are handled by the program, and by enhancing its error reporting capabilities.

Number: Patch 1214.00

Abstract: Fixes binlogd problem

State: Supersedes Patches 586.00, 777.00

This patch:

- Fixes a problem in which the binlog daemon can core dump if it attempts to recover events from a panic dump file containing invalid event data.
 - Fixes a time formatting problem when Compaq Analyze is used to display events in time zones with a positive offset from GMT.
 - Causes the binary error log daemon, binlogd, to sync its logfiles before closing them on system shutdown.
-

Number: Patch 1220.00

Abstract: Added support for DECthreads V3.18-150

State: Supersedes Patches 82.00, 420.00, 783.00, 785.00, 1218.00

This patch installs DECthreads V3.18-150, which is the latest support version of the HP POSIX Threads library for Tru64 UNIX V5.1A. It corrects several problems.

Number: Patch 1222.00

Abstract: Problem using MD5 authentication with Version 2 RIP

State: New

This patch corrects a problem using MD5 authentication with Version 2 RIP.

Number: Patch 1224.00

Abstract: Security

State: New

This patch fixes a potential security vulnerability which, under certain circumstances, could compromise system integrity. This may be in the form of improper file or privilege management.

Number: Patch 1226.00

Abstract: Security

State: New

This patch fixes a potential security vulnerability which, under certain circumstances, could compromise system integrity. This may be in the form of improper file access.

Number: Patch 1228.00

Abstract: Security

State: New

This patch fixes a potential security vulnerability which, under certain circumstances, could compromise system integrity. This may be in the form of improper file or privilege management.

Number: Patch 1226.00

Abstract: Security

State: New

This patch fixes a potential security vulnerability which, under certain circumstances, could compromise system integrity. This may be in the form of improper file or privilege management.

Number: Patch 1232.00

Abstract: Security (SSRT2368, SSRT2368)

State: New

This patch fixes a potential security vulnerability which, under certain circumstances, could compromise system integrity. This may be in the form of improper file or privilege management.

Number: Patch 1235.00

Abstract: SDLT media error caused bus resets with KZPCA adapter

State: Supersedes Patches 261.00,483.00, 797.00, 1233.00

This patch:

- Fixes a panic caused by SCSI bus resets with KZPCA HBAs.
 - Fixes a kernel memory fault panic after an "ITPSA: itpsa_action - error converting path ID to ITPSA softc structure" message.
 - Adds the capability for KZPCA devices to work with SCSI devices that only support asynchronous data transfers.
 - Fixes a kernel memory fault related to the KZPCA adapter.
 - Fixes SDLT media error caused bus resets with KZPCA adapters.
-

Number: Patch 1237.00

Abstract: Security (SSRT2193)

State: New

This patch fixes a potential security vulnerability which, under certain circumstances, could compromise system integrity when a buffer overflow occurs in the mailcv 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 1240.00

Abstract: Fixes for the collect utility

State: Supersedes Patches 175.00, 799.00, 1238.00

This patch installs Version 2.0.3 of the collect utility and does the following:

- Fixes a problem when collect is run in historical mode.
 - Fixes collect to correctly report the network interface load percentage.
 - Allows the collect to recognize and gather KZPCC disk statistics.
-

Number: Patch 1242.00

Abstract: libpset shared library fix

State: New

This patch fixes a problem that would prohibit processor set exclusive use on non-uniform memory access (NUMA) machines.

Number: Patch 1244.00

Abstract: libpset static library fix

State: New

This patch fixes a problem that would prohibit processor set exclusive use on non-uniform memory access (NUMA) machines.

Number: Patch 1246.00

Abstract: Security

State: New

This patch fixes a potential security vulnerability which, under certain circumstances, could compromise system integrity. This may be in the form of improper file access.

Number: Patch 1251.00

Abstract: Security (SSRT2280)

State: New (Supersedes Patch 1249.00)

This patch fixes several potential security vulnerabilities where under certain circumstances, system integrity may be compromised when a buffer overflow occurs in the dterm 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 1253.00

Abstract: Security

State: New

This patch fixes a potential security vulnerability which, under certain circumstances, could compromise system integrity. This may be in the form of improper file or privilege management.

Number: Patch 1255.00

Abstract: Security

State: New

This patch fixes a potential security vulnerability which, under certain circumstances, could compromise system integrity. This may be in the form of improper file or privilege management.

Number: Patch 1257.00

Abstract: Security

State: New

This patch fixes a potential security vulnerability which, under certain circumstances, could compromise system integrity. This may be in the form of improper file access.

Number: Patch 1260.00

Abstract: Fixes a problem in procfs

State: Supersedes Patches 216.00, 805.00, 1258.00

This patch:

- Fixes a kernel memory fault in procfs.mod.
 - Fixes a system panic from procfs ioctl user code.
 - Fixes a VM locking problem in procfs.
 - Fixes a kernel memory fault related to ioctl PIOCMAP.
 - Fixes a problem in procfs that, in some situations, prevents exiting threads from exiting. This creates a situation where these threads simply spin, consuming CPU time.
-

Number: Patch 1262.00

Abstract: Security

State: New

This patch fixes a potential security vulnerability which, under certain circumstances, could compromise system integrity. This may be in the form of improper file or privilege management.

Number: Patch 1264.00

Abstract: Security

State: New

This patch fixes a potential security vulnerability which, under certain circumstances, could compromise system integrity. This may be in the form of improper file access.

Number: Patch 1266.00

Abstract: Security

State: New

This patch fixes a potential security vulnerability which, under certain circumstances, could compromise system integrity. This may be in the form of improper file or privilege management.

Number: Patch 1270.00

Abstract: Fix for the hwmgr utility

State: Supersedes Patches 230.00, 811.00, 1267.00, 1268.00

This patch makes the following changes to the hwmgr utility:

- Corrects an error message that is displayed when hwmgr offlines a CPU that has only one bound process.
- Corrects the missing path to the SCP device the hwmgr -view devices command is issued.
- Lets hwmgr show CPU bindings with a tilde (~) character. Using the tilde will help distinguish between CPU bindings and RAD bindings, and will also keep the two interfaces consistent.
- Corrects a problem that causes hwmgr to display successful exit status for failed delete and unconfigure commands.
- Lets the hwmgr -view transaction -cluster command work on a cluster.
- Corrects some command-parsing irregularities in hwmgr that may cause options like -category and -cluster to be confused.

Number: Patch 1272.00

Abstract: Security

State: New

This patch fixes a potential security vulnerability which, under certain circumstances, could compromise system integrity. This may be in the form of improper file or privilege management.

Number: Patch 1274.00

Abstract: Security

State: New

This patch fixes a potential security vulnerability which, under certain circumstances, could compromise system integrity. This may be in the form of improper file or privilege management.

Number: Patch 1276.00

Abstract: Security (SSRT1-40U, SSRT1-41U, SSRT1-42U, SSRT1-45U)

State: Supersedes Patch 319.00

This patch fixes potential security vulnerabilities which, under certain circumstances, could compromise system integrity. These may be in the form of improper file access.

Number: Patch 1278.00

Abstract: Fix for bcheckrc script

State: New

This patch clarifies a vague message that dn_setup would print if a system was brought to single-user mode because dsfmgr detected database errors during boot.

Number: Patch 1281.00

Abstract: Corrects a problem with SNMP

State: New (Supersedes Patch 1279.00)

This patch:

- Corrects the problem where, when monitoring a system network, the output of IP Datagrams Received and IP Datagrams Sent looks strange when using Area graphs.
 - Corrects a problem with SNMP (Simple Network Management Protocol) retrieval of ARP (Address Resolution Protocol) cache data when an ARP table has more than 288 entries.
-

Number: Patch 1287.00

Abstract: Security (SSRT2280)

State: Supersedes Patches 561.00, 547.00, 813.00, 1284.00, 1285.00

This patch:

- Fixes a problem in the X Toolkit library (Xt) which could cause the TeMIP Iconic_map Presentation Module application (mcc_iconic_map) to crash.
- Fixes a problem where the definition of the X Toolkit function XtPending() was changed in Tru64 UNIX V5.1A, which caused some applications built on earlier versions of Tru64 UNIX to fail.
- Fixes a problem with Worldwide Language Support on Tru64 UNIX V5.1A where input method servers (such as VJE Delta) could not connect to their clients. This prevented users from entering non-English strings to their X Window System applications.
- Fixes several potential security vulnerabilities where under certain circumstances, system integrity may be compromised when a buffer overflow occurs in X11 applications. 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 1291.00

Abstract: Security (SSRT2280)

State: Supersedes Patches 549.00, 815.00, 563.00, 1288.00, 1289.00

This patch:

- Fixes a problem in the X Toolkit library (Xt) which could cause the TeMIP Iconic_map Presentation Module application (mcc_iconic_map) to crash.
- Fixes a problem where the definition of the X Toolkit function XtPending() was changed in Tru64 UNIX V5.1A, which caused some applications built on earlier versions of Tru64 UNIX to fail.
- Fixes a problem with cut and paste of JISX0212 Japanese characters on X Window System applications.
- Fixes a problem with Worldwide Language Support where input method servers (such as VJE Delta) could not connect to their clients, thereby preventing users from entering non-English strings to their X Window System applications.
- Fixes several potential security vulnerabilities where under certain circumstances, system integrity may be compromised when a buffer overflow occurs in X11 applications. 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 1293.00

Abstract: Security

State: New

This patch fixes a potential security vulnerability where under certain circumstances, system integrity may be compromised when a buffer overflow occurs in X11 applications. 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 1297.00

Abstract: Security

State: New

This patch fixes a potential security vulnerability which, under certain circumstances, could compromise system integrity. This may be in the form of improper file access.

Number: Patch 1299.00

Abstract: Security

State: New

This patch fixes a potential security vulnerability which, under certain circumstances, could compromise system integrity. This may be in the form of improper file access.

Number: Patch 1302.00

Abstract: Fix for fixfdmn problems

State: Supersedes Patches 177.00, 559.00, 818.00, 819.00, 820.00, 822.00, 1300.00

This patch makes the following changes to the fixfdmn utility:

- Fixes several problems where under extreme cases, it was possible for fixfdmn to core dump or to terminate without fixing the domain.
- Fixes a problem in which fixfdmn exits prematurely with the message "Can't allocate 0 bytes for group use array" and then instructs the user on how to make more memory available, even though more memory is not needed.
- Allows fixfdmn to modify only one page of the transaction log.
- Prevents fixfdmn from changing file sizes unnecessarily.
- Fixes a case where fixfdmn would abort when the same mcell was on the DDL more than once.
- Allows fixfdmn to be run on domains that have been mounted under Version 5.1B and then moved back to an older version.
- Fixes a problem in which fixfdmn could core dump on a rare corruption in the tag file.
- Allows fixfdmn to remove full frag groups from the free frag list in the fileset frag file.

Number: Patch 1305.00

Abstract: Security (SSRT0664U, SSRT0762U, SSRT0801)

State: Supersedes Patches 204.00, 206.00, 567.00, 1303.00

This patch:

- Fixes several potential security vulnerabilities which, under certain circumstances, could compromise system integrity. These may be in the form of improper file or privilege management.
- Corrects a problem with the ftpd daemon which could result in PC ftp clients hanging when transferring some files in ASCII mode.
- Prevents an ftp daemon failure when using globbing string of several asterisks and provides corrections for the help command and character drop with the put command.
- Corrects a problem in the ftp open command to allow the optional port argument to accept port numbers between 32768 and 65535.

Number: Patch 1307.00

Abstract: Fixes a problem with the operation of osf_boot

State: New

This patch fixes a problem with the operation of osf_boot.

Number: Patch 1309.00

Abstract: Security

State: New

This patch fixes a potential security vulnerability which, under certain circumstances, could compromise system integrity. This may be in the form of improper file or privilege management.

Number: Patch 1311.00

Abstract: Security

State: New

This patch fixes a potential security vulnerability which, under certain circumstances, could compromise system integrity. This may be in the form of improper file or privilege management.

Number: Patch 1313.00

Abstract: Security

State: New

This patch fixes a potential security vulnerability which, under certain circumstances, could compromise system integrity. This may be in the form of improper file or privilege management.

Number: Patch 1315.00

Abstract: Security (SSRT2191)

State: New

This patch fixes a potential security vulnerability where under certain circumstances, system integrity may be compromised when a buffer overflow occurs in the quot 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 1317.00

Abstract: Security (SSRT2191)

State: New

This patch fixes a potential security vulnerability where under certain circumstances, system integrity may be compromised when a buffer overflow occurs in the quot 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 1319.00

Abstract: Security

State: New

This patch fixes a potential security vulnerability which, under certain circumstances, could compromise system integrity. This may be in the form of improper file or privilege management.

Number: Patch 1328.00

Abstract: Security (SSRT0788U, SSRT0753U, SSRT0752U)

State: Supersedes Patches 493.00,495.00, 725.00

This patch:

- Fixes a potential security vulnerability which, under certain circumstances, could compromise system integrity. This may be in the form of large values of ENVIRONMENT variables.
 - Fixes 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.
 - Fixes a libXm.so incompatibility.
 - Fixes a problem with the Motif ToggleButton Widget where, in some cases, it may not draw itself correctly.
-

Number: Patch 1330.00

Abstract: Security (SSRT0788U, SSRT0753U, SSRT0752U)

State: Supersedes Patches 496.00, 498.00, 727.00

This patch:

- Fixes a libXm.so incompatibility.
- Fixes a problem with the Motif ToggleButton Widget where, in some cases, it may not draw itself correctly.
- Fixes 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.

A potential security vulnerability has been discovered, where under certain circumstances, system integrity may be compromised. This may be in the form of improper file or privilege management. HP has corrected this potential vulnerability.

Number: Patch 1332.00

Abstract: Security

State: New

This patch fixes 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 1334.00

Abstract: Security

State: Supersedes Patches 193.00 and 1119.00

This patch:

- Fixes several potential security vulnerabilities which, under certain circumstances, could compromise system integrity. This may be in the form of improper file access.
 - Improves the cleanPR script to clear Persistent Reservations on HSV110 device and to continue to go through all devices even if certain errors occur to one or some of devices.
 - Adds support in the cleanPR script to remove all Persistent Reservations for MSA controller.
-

Number: Patch 1336.00

Abstract: Security

State: New (Supersedes Patches 1193.00 and 1195.00)

This patch fixes several potential security vulnerabilities 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 1346.00

Abstract: Security

State: Supersedes Patches 67.00, 582.00, 574.00, 576.00, 703.00, 839.00, 1012.00, 1013.00, 1014.00, 1015.00, 1017.00, 1344.00

This patch:

- Prevents a vold from core dumping when removing a disk from rootdg using voldiskadm or voldg.
 - Prevents a KMF (kernel memory fault) panic, in voldiskstart(), when an I/O is attempted on an LSM device that is not accessible.
 - Prevents vold from core dumping on booting with the message:
commit: not holding dg lock <diskgroup>
 - Proves the enabler for Enterprise Volume Manager product.
 - Fixes a situation in which when a cluster member fails, mirrored volumes are left in a state such that recovery is always necessary when members boot, even if no additional recovery should be necessary.
 - Fixes a collision problem with clsm sync and LSM startup.
 - Fixes a problem of a vold core dump when old config db exists.
 - Fixes a problem where cluster node panics on boot if klog does not exist.
 - Fixes a problem of LSM not recognizing third-party disks.
 - Fixes an inability to create a new diskgroup when vold is in noloadbalance mode.
 - Fixes error messages for non-rootdg disks when cluster root is under LSM control.
 - Fixes problems in LSM's autoconfiguration feature, as well as some problems in the LSM commands volsave, volrestore, and volclonedg.
 - Fixes several potential security vulnerabilities where, under certain circumstances, system integrity may be compromised.
 - Allows the proper synchronization of disk errors and failures in a cluster under CLSM control.
 - Prevents inconsistent LSM volumes that can occur when the name of a partition that is being encapsulated matches the name of a current LSM volume.
 - Corrects a problem in which the SysMan Station daemon, smsd, was triggering LSM configuration errors when querying LSM in a cluster. This patch 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.
-

Number: Patch 1348.00

Abstract: Security (SSRT2260, SSRT1-40U, SSRT1-41U, SSRT1-42U)

State: Supersedes Patches 288.00, 507.00, 509.00, 1076.00, 1077.00, 1078.00, 1079.00, 1080.00, 1081.00, 1082.00, 1083.00, 1085.00

This patch:

- Fixes the following problems with the C shell command interpreter, csh:
 - Corrects the error message displayed when a non-root user performs issues the ls command with wild card characters on a directory having permission 700.
 - Corrects the error message displayed when nonomatch is set and a user performs issues the ls command with the question mark (?) character.
 - Fixes a problem so csh correctly recognizes the backslash (\) meta character.
 - Corrects the problem in which a user may experience a core dump when using csh from the Japanese locale.
- Fixes several potential security vulnerabilities where under certain circumstances, system integrity may be compromised when a buffer overflow occurs in the lpq, lpr and lprm commnads, the csh or dxterm utilities, or the telnetd daemon. Buffer overflows are sometimes exploited in an attempt to subvert the function of a privileged program and possibly execute commnads and the elevated privileges if the program file has the setuid privilege.
- Fixes several potential security vulnerabilities where, under certain circumstances, system integrity may be compromised. These may be in the form of improper file access.
- Corrects an lpc regression in the lpc buffer overflow fix.
- Corrects a problem where telnetd leaves an extra udp port open.
- Fixes a problem with csh redirection that can occur while redirecting standard input and standard output of a command to a file exist in home directory using the tilde (~) operation.

Number: Patch 1350.00

Abstract: Security (SSRT2193)

State: New (Supersedes Patches 1150.00, 1230.00)

This patch fixes 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 1352.00

Abstract: Support for Ultrium tape drive

State: New (Supersedes Patches 1159.00, 1160.00, 1161.00, 1162.00, 1164.00)

This patch:

- Adds support for new devices devices branded by HP.
 - Adds latent device recognition support for MSA1000 storage array controllers.
 - Enhances SuperDLT maximum transfer size edit to be more tolerant of previous changes.
 - Provides device support for the SDLT160/320 tape drive.
 - Adds support for Ultrium 2 SCSI tape drive.
 - Adds latent support for logical devices.
 - Provides support for the Ultrium tape drive.
-

Number: Patch 1357.00

Abstract: Security (SSRT0785U)

State: Supersedes Patches 472.00, 473.00, 474.00, 475.00, 477.00, 721.00, 723.00, 1066.00, 1067.00, 1068.00, 1069.00, 1071.00

This patch:

- Corrects a problem with dxaccounts in which a core dump occurs when /etc/shells is a directory instead of a file.
- Corrects a problem with dxaccounts in which the hour glass cursor remains after a failure to create a home directory in the process of adding or modifying an account.
- Fixes a problem of dxaccounts in which names and security attributes of UNIX users are not mapped correctly when they are viewed from PC Users' dialog.
- Fixes the problem that user name entries are replicated in the /etc/group file when modifying users with either dxaccounts or sysman accounts.
- Fixes a problem in dxaccounts that can cause certain C2 security values to not be displayed, which could result in unexpected values being saved.
- Fixes the problem of useradd, usermod, and dxaccounts ignoring password length restrictions when changing passwords.
- Fixes a number of problems with Dxaccounts on a system with ASU installed.
- Fixes several minor problems with the account management tools.
- Corrects a problem with the userdel command core dumping when the shell field is empty in the passwd file.
- Corrects a problem of the usermod command not working as expected with NIS +/- users.
- Corrects the problem where the useradd command fails to create a user with the specified template field under Enhanced Security.
- 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 several potential security vulnerabilities where, under certain circumstances, system integrity may be compromised.
- Corrects a condition in which the useradd command was not managing default and template data properly. This showed up most notably with useradd -p producing: "Password must be between 32 and 80 characters."

Number: Patch 1359.00

Abstract: Corrects hang in the log command

State: Existing

This patch corrects a possible deadlock in the ./isl/log and ./usr/sbin/log commands.

Number: Patch 1367.00

Abstract: Security (SSRT1-40U, SSRT1-41U, SSRT1-42U, SSRT1-45U)

State: Supersedes Patches 181.00, 486.00, 488.00, 191.00, 2.00, 121.00, 241.00, 243.00, 400.00, 401.00, 402.00, 403.00, 404.00, 405.00, 406.00, 407.00, 408.00, 409.00, 410.00, 412.00, 307.00, 302.00, 162.00, 253.00, 255.00, 90.00, 218.00, 303.00, 305.00, 421.00, 422.00, 423.00, 424.00, 426.00, 681.00, 682.00, 683.00, 684.00, 685.00, 686.00, 687.00, 689.00, 807.00, 511.00, 729.00, 541.00, 146.00, 148.00, 126.00, 127.00, 128.00, 129.00, 130.00, 131.00, 132.00, 134.00, 286.00, 6.00, 7.00, 8.00, 9.00, 10.00, 11.00, 12.00, 13.00, 14.00, 15.00, 16.00, 17.00, 18.00, 19.00, 20.00, 21.00, 22.00, 23.00, 24.00, 25.00, 26.00, 27.00, 28.00, 29.00, 30.00, 31.00, 32.00, 33.00, 34.00, 35.00, 36.00, 37.00, 38.00, 39.00, 40.00, 41.00, 42.00, 43.00, 44.00, 45.00, 46.00, 47.00, 48.00, 49.00, 50.00, 51.00, 52.00, 53.00, 54.00, 55.00, 56.00, 57.00, 58.00, 59.00, 60.00, 61.00, 102.00, 63.00, 104.00, 214.00, 236.00, 246.00, 247.00, 248.00, 250.00, 270.00, 271.00, 272.00, 273.00, 274.00, 275.00, 276.00, 277.00, 279.00, 296.00, 298.00, 321.00, 323.00, 325.00, 290.00, 152.00, 93.00, 95.00, 328.00, 329.00, 330.00, 331.00, 332.00, 333.00, 334.00, 335.00, 336.00, 337.00, 338.00, 339.00, 340.00, 341.00, 342.00, 343.00, 344.00, 345.00, 346.00, 347.00, 348.00, 349.00, 350.00, 351.00, 352.00, 353.00, 354.00, 355.00, 356.00, 357.00, 358.00, 359.00, 360.00, 361.00) 362.00, 363.00, 364.00, 365.00, 366.00, 367.00, 368.00, 369.00, 370.00, 371.00, 372.00, 373.00, 374.00, 375.00, 376.00, 377.00, 378.00, 379.00, 380.00, 381.00, 382.00, 383.00, 384.00, 385.00, 386.00, 387.00, 388.00, 389.00, 390.00, 391.00, 392.00, 393.00, 394.00, 395.00, 396.00, 397.00, 399.00, 543.00, 590.00, 592.00, 596.00, 203.00, 594.00, 597.00, 598.00, 599.00, 600.00, 601.00, 602.00, 603.00, 604.00, 605.00, 606.00, 607.00, 608.00, 609.00, 610.00, 611.00, 612.00, 613.00, 614.00, 615.00, 616.00, 617.00, 618.00, 619.00, 620.00, 621.00, 622.00, 623.00, 624.00, 625.00, 626.00, 627.00, 628.00, 629.00, 630.00, 631.00, 632.00, 633.00, 634.00, 635.00, 636.00, 637.00, 638.00, 639.00, 640.00, 641.00, 642.00, 643.00, 644.00, 645.00, 353 646.00, 647.00, 648.00, 649.00, 650.00, 651.00, 652.00, 653.00, 654.00, 655.00, 656.00, 657.00, 658.00, 659.00, 660.00, 661.00, 662.00, 663.00, 664.00, 665.00, 666.00, 667.00, 668.00, 669.00, 671.00, 672.00, 673.00, 674.00, 675.00, 676.00, 677.00, 678.00, 680.00, 834.00, 835.00, 837.00, 842.00, 846.00, 853.00, 854.00, 855.00, 856.00, 857.00, 858.00, 859.00, 860.00, 861.00, 862.00, 863.00, (864.00, 865.00, 866.00, 867.00, 868.00, 869.00, 870.00, 871.00, 872.00, 873.00, 874.00, 875.00, 876.00, 877.00, 878.00, 879.00, 880.00, 881.00, 882.00, 883.00, 884.00, 885.00, 886.00, 887.00, 888.00, 889.00, 890.00, 891.00, 892.00, 893.00, 894.00, 895.00, 896.00, 897.00, 898.00, 899.00, 900.00, 901.00, 902.00, 903.00, 904.00, 905.00, 906.00, 907.00, 908.00, 909.00, 910.00, 911.00, 912.00, 913.00, 914.00, (915.00, 916.00, 917.00, 918.00, 919.00, 920.00, 921.00, 922.00, 923.00, 924.00, 925.00, 926.00, 652 927.00, 928.00, 929.00, 930.00, 931.00, 932.00, 933.00, 934.00, 935.00, 936.00, 937.00, 938.00, 939.00, 940.00, 941.00, 942.00, 943.00, 944.00, 945.00, 946.00, 947.00, 948.00, 949.00, 950.00, 951.00, 952.00, 953.00, 954.00, 955.00, 956.00, 957.00, 958.00, 959.00, 960.00, 961.00, 962.00, 963.00, 964.00, 965.00, 966.00, 967.00, 968.00, 969.00, 970.00, 971.00, 972.00, 973.00, 974.00, 975.00, 976.00, 977.00, 978.00, 979.00, 980.00, 981.00, 982.00, 983.00, 984.00, 985.00, 986.00, 987.00, 988.00, 80.00, 418.00, 780.00, 782.00, 1215.00, 990.00, 1217.00, 1248.00, 1320.00, 1322.00, 1323.00, 1324.00, 1326.00, 1338.00, 1339.00, 1340.00, 1341.00, 1343.00, 1355.00, 1360.00, 1362.00, 1363.00, 1364.00, 1365.00

This patch:

- Fixes several potential security vulnerabilities where, under certain circumstances, system integrity may be compromised. These may be in the form of improper file access.
 - Fixes 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 a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of file corruption due to the manner in which setuid/setgid programs core dump.
 - Fixes a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of network programs core dumping.
-

Patch 1367.00 Continued

- Fixes a potential security vulnerability where, under certain circumstances, users can clobber temporary files created by shell commands and utilities; for example, under /sbin, /usr/sbin, /usr/bin, and /etc).
- Fixes a potential security vulnerability where under certain circumstances, system integrity may be compromised. This could result in a panic with the string: "lock_clear_recursive: recursion not enabled."
- Fixes a potential security vulnerability where, under certain circumstances, a race condition can occur that could allow a non-root user to modify any file and possibly gain root access.
- Fixes a potential security vulnerability where, under certain circumstances, a remote system can take over packets destined for another host.
- Fixes a potential security vulnerability where, under certain circumstances, system integrity may be compromised when a buffer overflow occurs in the BIND 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.
- Fixes a potential security vulnerability, where under certain circumstances, system integrity may be compromised when a buffer overflow occurs in the xdr library, which is used by the rpc library. 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.
- Fixes a potential security vulnerability where, under certain circumstances, system integrity may be compromised when a buffer overflow occurs in the chfn, chsh, or passwd utilities. 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.
- Fixes a potential security vulnerability where, under certain circumstances, system integrity may be compromised when a buffer overflow occurs in the ypmatch and traceroute utilities. 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.
- Fixes a potential security vulnerability where, under certain circumstances, system integrity may be compromised when a buffer overflow occurs. 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.
- 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)

- 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:

SSRT2322 - BIND resolver (Severity - High)

- Fixes a class scheduler semaphore race condition. The class scheduler depends on semaphores to protect its database from simultaneous updates. This patch automatically detects if the semaphore no longer exists and allocates a new one, allowing the class scheduler to proceed without interruption.
-

Patch 1367.00 Continued

- Fixes a problem where logins appear to be hung on standalone systems with Enhanced Security enabled.
 - Fixes a regular expression matching problem in multibyte locales.
 - Fixes the `-ignore_all_versions` and `-ignore_version` flags for the run-time loader (`/sbin/loader`).
 - Fixes a problem where `strtod()` was returning different outputs for the same input.
 - Fixes a problem where the `tan()` function was returning the wrong results.
 - Eliminates a libc memory leak that occurred when calling `dlclose()` in applications linked with the threads run-time environment.
 - Changes the optional dynamic loader arguments `-allocator_range` and `-allocator` to `-preallocated_range`.
 - Fixes a problem in `mktime()` when adjusting for a `tm` struct containing an invalid `tm_isdst` (daylight savings time) setting.
 - Fixes a segmentation fault problem with long `LOCPATH` and `LANG` values.
 - Fixes a problem in which the RPC TCP server incorrectly tries to write to a socket that has already been closed by a client.
 - Fixes an application core dump problem when the `LANG` environment variable is too long.
 - Fixes a problem with the `fopen()` function in which `fopen()` returned a "file not found" message when insufficient memory was available to allocate the `FILE` structure. With this patch, the `fopen()` function now returns the message "not enough space" for this case.
 - Fixes a problem in `fread()` in which excessive I/O was taking place for large amounts of data, causing performance problems. It also addresses a failure in `fread()` to properly handle data sizes that have representations greater than 32 bits (2^{32} of data).
 - Fixes a loader core dump that occurs when invoking certain `call_shared` executables that have been processed by post-link instrumentation tools.
 - Fixes a problem with the `strerror()` function in which buffers could not be allocated.
 - Fixes a problem with the `fwrite()` function in which it failed when the total number of bytes to be written is larger than 2 GB.
 - Fixes a regular expression problem with the `REG_NEWLINE` flag of the `regexexec()` routine.
 - Resolves a memory leak and a filtering issue in the Event Manager (EVM), and allows the `evmwatch` utility to reconnect automatically if `evmd` fails and is restarted.
 - Provides enablers for the Compaq Database Utility.
 - Fixes a problem in which binary error log (`binlog`) events posted by the EMX fibrechannel driver and the system console are reported incorrectly by the Event Manager.
 - Provides the `/usr/sbin/mkstemp` program which allows the mechanism to create a secure temporary file.
 - Fixes a problem in which the EVM daemon acting as a subscribing client within a cluster will unexpectedly drop the connection to the other EVM daemons in the cluster. This may happen when an EVM client subscribes to events specifying the cluster alias.
 - Resolves an issue which can cause an Event Manager client or the EVM daemon to core dump under rare circumstances.
 - Fixes a `sys_check` problem in which verification of invoking a processes' name in `CLISCRIP`T failed due to the `PARSING` of `ps` command output.
-

Patch 1367.00 Continued

- Fixes a multi-thread timing window in malloc and free where the list of free chunks could become corrupted, resulting in a segfault.
 - Fixes a regular expression performance problem in sed.
 - Fixes a problem with printing long double values.
 - Fixes a performance degradation in malloc, in applications which perform many mallocs and few frees. With this patch, the performance of malloc is constant regardless of the number of allocated chunks outstanding.
 - Fixes a problem with atexit() or pthread_atfork() handlers in shared libraries. An application will crash when handlers in shared libraries are called after the libraries are dlclosed and unmapped.
 - Fixes a problem in which compiled format doprnt code does not handle precision correctly.
 - Corrects the following problems with the alt driver for DEGPA Gigabit Ethernet adapters. These problems affect all Tru64 systems using alt with vMAC or NetRAIN:
 - Fixes vMAC to let it work with DEGPA.
 - Prevents two DEGPA adapters from getting the same MAC address in a NetRAIN configuration.
 - Fixes a problem with RLIMIT_DATA process limits when running fsck on a large file system.
 - Fixes "ata_probe: reset failed, sts=0x7f, err=0x7f" errors for IDE disks not connected to the system.
 - Fixes a rare emx driver boot issue and a rare heavy load I/O hang.
 - Updates the emx driver to v2.03 to fix a problem that could cause the driver panic during adapter resets.
 - A previous release of this patch updated the emx driver to v2.02 to correct the following problems:
 - A panic of cannot grow probe list
 - A problem of a mcs_lock panic when an adapter experiences a h/w hang condition
 - A previous release of this patch updated the emx driver to v2.01 to correct the following problems:
 - A problem of unexpected tape I/O aborts
 - Panic of “can’t grow probe list”
 - Several kernel memory faults within the driver
 - Redundant adapter failures no longer panic the system
 - A problem of panicing with low memory resources
 - Stalling I/O during reprobing when a cluster member goes down
 - Fixes problems seen with the loading and unloading of dynamic drivers.
 - Fixes a problem when using the VX1 graphics module in which the mouse cursor disappears when moved along left- and top-most edge.
 - Fixes a kernel crash dump generation problem that resulted in the wrong pages being compressed and written. Without this fix, postmortem debugging may be difficult or impossible.
 - Fixes a "simple_lock timeout" system panic due to a bug between mcs_unlock and mcs_lock_try on the same CPU.
-

Patch 1367.00 Continued

- Provides New Hardware Delivery V4 (NHD4) enablers for future hardware support including:
 - Graphic devices
 - A new platform
 - An array controller
 - Fixes a domain panic pointing to quotaUndo, when a domain has a fileset with a clone, the clone is deleting, and a file in the fileset finds no space available in the domain.
 - Provides a new /usr/sbin/wol command that utilizes the Wake feature for a future platform.
 - Fixes a time loss problem seen on DS systems only when using console callbacks. The patch resynchronizes the clock when time loss is detected.
 - Fixes a rare panic in the driver for the DE600/DE602 10/100 Ethernet adapter.
 - Corrects a problem where the network subsystem sometimes sends a null TCP packet when a connection is reset.
 - Provides enabler support for Enterprise Volume Manager product.
 - Fixes a system panic with "malloc_check_checksum: memory pool corruption."
 - Fixes a problem in which issuing a quot -h command causes a memory fault when the /etc/fstab file contains a mount point that is not mounted.
 - Fixes a problem with IPv6 raw socket creations.
 - Corrects a CFS problem that could cause a panic with the panic string of "CFS_INFS full."
 - Fixes a problem with erroneous data being returned from the DEVIOCGET ioctl if an error occurs while processing the ioctl.
 - Fixes a problem in which a TCP socket can continue to receive data with no application running.
 - Fixes a performance problem and the results are large performance increases in configurations where more than eight tapes are supported on a Fibre Channel (usually behind an MDR or FCTCII).
 - Allows a single ddr.dbase entry to support a particular SCSI device on both parallel SCSI and FC busses. Previously, SCSI devices connected behind an FCTCII or MDR would not be properly associated with their ddr.dbase entry.
 - Fixes problems found with RAID Services that include:
 - RAID services not acknowledging presence of CAM RAID device
 - A hang, the inability to prohibit a user from deleting a logical volume while it is in use
 - A "malloc_check_checksum: memory pool corruption" system panic
 - Fixes a problem in which threads can hang in x_load_inmem_xtnt_map().
 - Fixes the following Virtual Memory problems. The first three are seen on NUMA systems only, and the fourth problem can be seen on any system type:
 - A "vm_pg_alloc: page not free" system panic that occurs during process migration.
 - A "vm_pageout_activate: page already active" system panic that occurs if one thread is unlocking some pages in memory while another thread is migrating them.
 - Memory inconsistencies caused by a fault path for large shared memory regions prematurely releasing a hold on a page it just locked. This can cause variety of problems including user program errors and system panics.
 - A "simple_lock: time limit exceeded" system panic that occurs if very large (8MB or larger) System V Shared memory regions are in use.
-

Patch 1367.00 Continued

- Fixes a problem in which a panic occurs while task swapping.
 - Fixes a problem in virtual memory that can cause a kernel memory fault.
 - Fixes a problem in which the I/O transfer rate can suddenly drop when writing to a hole in an AdvFS domain, when a volume in that domain becomes full.
 - Fixes a problem with the memory troller attempting to post an EVM event indicating that a particular PFN has been mapped out.
 - Fixes lock time issues and UBC performance problems, and provides AdvFS and UFS performance improvements in systems (other than the AlphaServer GS80, GS160, and GS320) with low memory.
 - Fixes several problems related to shared memory (memory that can be accessed by more than one CPU) that could lead to panics, hangs, and performance problems.
 - Fixes a bug that can cause performance problems for certain applications when the sysconfigtab parameter ipc:sem_broadcast_wakeup is set to 0.
 - Fixes a problem in which a check for managed address may return an invalid value when called with the address of a gh region not on rad 0.
 - Fixes a kernel memory fault in msg_rpc_trap.
 - Fixes a potential problem with lost data after a direct I/O write with a file extension followed quickly by a system crash.
 - Fixes a crash that occurs when disk controllers are restarted repeatedly.
 - Fixes a "u_shm_oop_deallocate: reference count mismatch" error due to a problem in the locking mechanism when gh_chunks are in use.
 - Provides I/O barrier code that prevents HSG80 controller crashes (firmware issue).
 - Corrects the problem of a thread deadlocking against itself under the following conditions:
 - Running in a cluster.
 - Opening (and then closing) a directory that has an index file.
 - Trying to open the index file through .tags (for example, defragment) and by coincidence getting the vnode that pointed to the directory that the index file is attached to.
 - Fixes kernel panic "bs_invalidate_rsvd_access_struct: bad access struct."
 - Ensures that DMAPI region information maintains consistency across CFS server and client nodes in the case that an unexpected node failure occurs.
 - Fixes a problem where additional HSZ70 control ports, /dev/cport/scpN, were created during HSZ70 controller failover operations.
 - Prevents a crash seen while deleting SCSI devices using hwmgr.
 - Fixes a problem where new devices could be created when following the HSZ70 controller failover procedure.
 - Fixes a problem in which reading a clone file that is still in the cache after using the rmvol utility may panic the system.
 - Fixes a problem where a variable was used without being initialized, which could lead to a possible kernel memory fault.
 - Provides AdvFs and VFS support for the freezefs and thawfs commands.
 - Provides EVM V2 enablers.
 - Provides several changes to CAM:
 - Fixes a problem where passthrough IOCTL fails with EIO (CAM_BUSY)
 - Fixes a RESERVATION CONFLICT driver BUSY problem
 - Enforces superuser-only access for SCSI passthrough.
-

Patch 1367.00 Continued

- Enables access to SCSI control ports (/dev/cport/scp??), allowing management of some types of RAID controllers.
 - Eliminates unintended AutoFS auto-mount storms.
 - Fixes a problem where extraneous "This node removed from cluster" events cause panics of cluster nodes.
 - Fixes a panic that occurs if DMAPI operations are erroneously executed on an NFS file system.
 - Processes triggering stack growth with anon_rss_enforce set to 2, and exceeding the set resident memory limit hang or panic.
 - Fixes a kernel panic with "xfer_hole_stg: unaligned kernel access" or "xfer_hole_stg: kernel memory fault."
 - Fixes a timing window where flushing data to disk can be incomplete when a system is going down, if more than one thread calls reboot() without first going through shutdown, /sbin/reboot, or /sbin/halt.
 - Ensures that if an AdvFS file is opened for both O_DIRECTIO and O_APPEND, threads racing to append data to the file will be correctly synchronized, and all data will be appended to the file.
 - Fixes several direct I/O problems seen when using the aio interface. The symptoms include a kernel memory fault, and an aio condition that causes a live_dump to be generated.
 - Fixes a condition where the smoothsync thread, in attempting to flush dirty buffers for memory-mapped files, would also flush buffers for non-memory-mapped files. This did not cause any errors, but could cause more I/O than necessary to be done.
 - Allows POSIX semaphores/msg queues to operate properly on a CFS client.
 - Fixes a problem where running verify may panic the system.
 - Corrects a condition in which a kernel memory fault may occur while attempting to read a log record.
 - Prevents a race in msfs_umount.
 - Provides a fix to a deadlock situation that can occur when invoking the hwmgr -show comp command while the devices on an HSZ70 are changing their names.
 - Fixes a problem where network interfaces can appear unresponsive to network traffic.
 - Fixes a problem where "path reduced" messages are printed at boot time for devices that still have at least one valid path.
 - Enables the quick reclaim and deallocation of a vnode.
 - Fixes a problem where a panic may occur when the DMAPI functionality is in use.
 - Fixes a problem where the setgid bit of a directory was not being set when created, if its parent directory had the setgid bit set.
 - Makes several changes to kernel routing:
 - Fixes a problem that caused a panic when deleting an IP address.
 - Fixes a problem of a panic when performing IP reconfiguration.
 - Adds an interface route on address configuration.
 - Fixes a problem that caused an "ics_unable_to_make_progress: input thread stalled" panic.
-

Patch 1367.00 Continued

- Addresses three UBC issues:
 - Reinstates `abc_maxpercent` hardlimit behavior.
 - Allows the UBC to purge and steal pages under very low free memory conditions during page allocation.
 - Removes memory mapping for NFS pages being invalidated and freed. Pages were being freed but still mapped to the process.
 - Fixes NFS support for the Enterprise Volume Manager product.
 - Corrects a performance problem in which NFS V3 I/O used larger than necessary buffers when writing to locked files, resulting in lower throughput.
 - Provides a script, `/usr/sbin/evm_versw_undo`, that will allow a user to remove the EVM patch after the version switch has been thrown by running `clu_upgrade -switch`. This script will set back the version identifiers and request a cluster shutdown and reboot to finish the deletion of the patch. Another rolling upgrade will be required to delete the patch with `dupatch`.
 - Enables a version-switched patch.
 - Causes a SCSI check condition with NO SENSE status to be treated by the disk driver as a condition to retry the I/O.
 - Fixes a condition in which a panic that could occur if an illegal argument is passed to UFS mount by a root user.
 - Fixes a kernel build failure when AdvFS is excluded from the build.
 - Fixes a problem where the system may be hung or there are poor response times on systems with limited numbers of CPUs.
 - Fixes a condition that causes an "RDG unwire panic" when running with RDG and GH chunks.
 - Resolves a problem where duplicate attributes are registered for all CAM devices present in a system. This affects `iostat` output and any other application that relies on the attribute data.
 - Adds workarounds for additional firmware problems found in the HSx controller.
 - Fixes for scheduler at high load averages and initial NUMA process placement.
 - Fixes an `rmvol` failure that would be seen as an `E_PAGE_NOT_MAPPED` error when no more space is available for user data migration to another volume in the domain.
 - Fixes the following tape drive problems:
 - Tape devices in multipath configurations unexpectedly rewind or go off line. (Multipath means that I/O can reach the device by an alternate data path, such as a redundant controller or bus.) Note that this patch reverts the tape drive configuration to single path mode.
 - The `vdump` utility fails to close because the drive goes offline before the dump operation is complete. An error message similar to the following is displayed:

```
vdump: unable to properly close device </dev/tape/tape1_d1>; [5] I/O error
```
 - Fixes a problem in which opening a disk partition sometimes fails when the disk is on shared bus.
 - Fixes a kernel memory fault panic on NUMA systems because of corrupt UBC LRU.
 - Fixes a problem of poor interactive response, including hanging commands and logins and random drops in I/O rates when writing many large files.
 - Fixes a potential problem in which stale data may be returned to an application running on a CFS client when it reads data from a file on a CFS server. Another possible symptom is incomplete flushing of user data when an `fsync()` is issued or an `O_[D]SYNC` write is performed.
-

Patch 1367.00 Continued

- Fixes a problem where new barrier code will not reserve after a registration if a new device or cluster is installed.
- Fixes a problem where HSV110 Persistent Reserve with a Reservation conflict SCSI status gets passed off to cam_notify when it should not, resulting in incorrect reservation status.
- Fixes a problem with data inconsistency that can occur when a CFS client reads a file that was recently written to.
- Fixes SEL logging problem where panic events were logged as misc events. It also adds new event types that can be logged.
- Fixes a problem in which the system could panic while performing CPU hotswap.
- Fixes a problem in which Link Aggregation groups can be successfully created and configured but are unable to successfully transmit and receive packets over the resulting lag interface.
- Prevents a potential panic with non-StorageWorks RAID controllers that used the same name for a controller and a disk drive. Although this conflict was resolved in a prior release, it was possible that an attempt by the kernel to access the disk drive could result in a system panic.
- Provides support for a related cluster patch.
- Removes a panic seen at boot time of the form:
panic (cpu 6): u_anon_oop_deallocate: anon_rss_pagelist has pages queued
- Fixes a kernel memory fault in wait_to_readyq(), advfs_page_busy(), or potentially other routines that may reference a vm_page, bsBuf or ioDesc that has been freed prematurely.
- Fixes the C++ incompatibility of the following:

```
/usr/include/io/dec/bi/bdareg.h
/usr/include/io/dec/bi/buareg.h
/usr/include/io/dec/eisa/aceregs.h
/usr/include/io/dec/eisa/eisa.h
/usr/include/io/dec/fbus/fbusreg.h
/usr/include/io/dec/pci/pci.h
/usr/include/io/dec/pcmcia/pcmcia.h
/usr/include/io/dec/pcmcia/ti1130_reg.h
/usr/include/io/dec/tc/sccreg.h
/usr/include/io/dec/tc/tc.h
/usr/include/io/dec/ws/comet_driver.h
/usr/include/io/dec/ws/comet_regs.h
/usr/include/io/dec/ws/inputdriver.h
/usr/include/io/dec/ws/ws_driver.h
```

- Restores the cam_logger() interface to its published specifications, and introduces the cam_logger3() interface to accept a hardware ID in its parameter list.
 - Addresses a potential UBC panic that could occur when accessing CFS file systems.
 - Fixes a problem with vm_faults against anon objects mapped by multiple map entries.
 - Provides ECC Enhancements for DTAG error logging for AlphaServer GS80, GS160, and GS320 systems.
 - Fixes a problem where decreasing the smoothsync_age does not always have an effect.
 - Fixes system panic and data corruption caused by changing the fifo parameter pipe-databuf-size while fifo operations are in flight.
 - Fixes AdvFS synchronization problems with lingering I/O messages during domain deactivation or rmvol actions.
-

Patch 1367.00 Continued

- Fixes problems caused by certain `kmem_debug` settings (`kmem_debug=0x40`, `kmem_protected_size=4096`) and AdvFS's handling of freed memory.
- Fixes and provides enhancements to Tru64 UNIX to support Encore realtime software.
- Modifies the `rmvol` command error messages to reflect why `rmvol` fails.
- Modifies the `showfdmn` command so it does not print success message on a failure. For example:

```
showfdmn: unable to get info for domain 'domain_used'  
showfdmn: Successful
```

- Fixes a potential CFS deadlock.
 - Fixes a problem when running `ssh v2.4.0` and `v2.4.1` when executing `ls` in `sftp` and when uploading public key using `ssh-pubkeymgr`.
 - Fixes SEL logging problem where panic events were logged as misc events. It also adds new event types that can be logged.
 - Corrects a problem that is encountered when trying to create an Oracle database on an AlphaServer GS80, GS160, or GS320 system that has a memoryless QBB. Without this patch, direct I/O to an AdvFS file using asynchronous I/O will hang if it is completed on a memoryless QBB.
 - Corrects problems when running the `dd` utility on a disk with a label. It would not return errors when expected.
 - Fixes a problem with I/O suspended (hung) in a cluster configuration where one or more rad does not have a valid, initialized path.
 - Fixes a problem that causes bugchecks from applications running `DecThreads`.
 - Fixes a problem for locking on retry case for multi-threaded `select/poll`. For example: `PANIC: "thread_block: simple lock owned."`
 - Fixes a potential problem where system responsiveness may be impacted.
 - Fixes a kernel memory fault in `DMAPI` code under cluster stress conditions.
 - Fixes a calculation leading to poor hash table distribution for NFS client mountpoints in the cluster.
 - Eliminates unintended AutoFS auto-mounts, in particular those that may result via the execution of the `df` command on systems running Tru64 UNIX versions earlier than Version 5.0.
 - Corrects a problem in which multi-volume AdvFS V3 domains exhibit I/O errors that are not attributable to hardware. The same problem also causes a failed `mkfs` due to `ENO_XTNTS`.
 - Corrects a problem with storage allocation by reducing the number of extent maps generated, subsequently giving better I/O performance on the resulting file. Prior to this modification, storage allocation for a file opened for direct I/O could, depending on the write sizes requested, have large extent maps even though the disk is not fragmented. Although the file functions correctly, performance is reduced by the numerous extent maps.
 - Corrects a problem in which file permissions inherited from a default ACL may be different than expected in rare cases.
 - Corrects a problem where the DLI queue stalls when there is no traffic in the TCP/IP or HDLC stacks.
 - Corrects a problem whereby clocks on systems could move backwards after subsequent relocations of the root file system using the `cfsmgr` command.
 - Corrects a problem where a "kernel stack not valid" halt on a CPU will trigger a "PANIC TB_SHOOT ACK TIMEOUT" or lock timeout on a non-NUMA system.
-

Patch 1367.00 Continued

- Corrects a problem with a simple lock timeout, or a panic due to holding a simple lock during a context switch on a non-NUMA system.
- Corrects an issue seen on NFS clients. The aggressive behaviour of client negative lookup cache for concurrent create/lookup was corrected.
- Corrects an issue with mmaped() files on a NFS-mounted file system in which changes to an mmaped() file were not immediately seen.
- Fixes a problem where the tape changer is only accessible from member that is the drd server for the changer.
- Fixes a problem where socket-based applications can hang in soclose().
- Fixes a problem during file system relocation in which the system may panic due to a kernel memory fault when a directory larger than 8192 bytes has been deleted, while simultaneously being accessed by another thread.
- Corrects a kernel memory fault on multiple CPU systems when two or more CPUs find an AdvFS problem at the same time.
- Fixes a problem where, after a system crash, on reboot there is a domain panic.
- Corrects the problem where attempts to delete psets can hang the system.
- Prevents an AdvFS metadata inconsistency in the event of a system crash.
- Prevents a possible extent map corruption when multiple volumes are full.
- Fixes a problem with multi-threaded applications that can cause the application to consume 100 percent of the CPU usage time.
- Fixes a domain panic in a cluster when a file system is mounted on a disk accessed remotely over the cluster interconnect.
- Fixes a locking problems in vclean().
- Fixes the CEH bus/target and lun number when lun > 127.
- Fixes a kernel memory fault when freeing devices.
- Corrects problems with USB causing panics on heavily stressed systems.
- Corrects a problem with the counters maintained for the NetRAIN virtual interface.
- Installs Version 1.02 of the Ciss Driver.
- Fixes a problem in which the psrinfo -v command may print an incorrect CPU cache size in a mixed CPU size/speed environment.
- Prevents a panic in assert_wait_mesg caused by the posting of an event_wait without clearing a previous request.
- Fixes a problem where tape and changer devices on fibre could occasionally return an incorrect offline status.
- Enables improvements in the kernel crash dump subsystem that make it possible to generate a dump after disk driver shutdown has taken place.
- Fixes a potential "kernel memory fault" panic in the Virtual Memory subsystem on SMP systems.
- Adds hardware support for the AlphaServer DS25.
- Fixes a minor problem in the AlphaServer ES45 environmental error handling code.
- Corrects problems where NFS can deadlock.
- Corrects an AdvFS problem in which EIOs are returned by AdvFS to NFS.
- Corrects a kernel memory fault panic in malloc_thread().
- Fixes the predictable TCP Sequence Number.
- Corrects a data inconsistency that can occur when a CFS client reads a file that was recently written to.

Patch 1367.00 Continued

- Provides support for a related cluster patch to support multiple filesets being mounted from the cluster_root domain.
 - Fixes a potential deadlock situation when using freezefs on multiple domains while also running addvol (or rmvol).
 - Fixes numerous problems of accessing de-allocated and freed vnodes.
 - Fixes the following problems in Link Aggregation (LAG):
 - An inability to modify the ipmtu of a LAG interface
 - An inability to work with gigabit Ethernet jumbo frames
 - Attempts to use a link that is down
 - Poor performance in server-to-server configurations
 - Fixes a situation which a failed open to a device causes an error that the device cannot be deleted using the hwmgr command.
 - Fixes an incorrect return type in a logging routine that prevented proper operation of the memory troller on an AlphaServer DS20L.
 - Corrects a problem when offlining a processor in which a seldom-taken code path may attempt to take a complex (sleep or blocking) lock while in interrupt context, thereby causing the kernel to panic.
 - Fixes a potential problem with vdf and showfdmn, where they could incorrectly display the message “showfdmn: No such file or directory.”
 - Prevents a cluster file system server panic that can occur if a cluster client clears the server cache entries for a file being operated on by defragment, balance, migrate, rmvol, or mssh.
 - Fixes several problems found in the KZPEA driver that could result in hung I/O, pending I/O not being cleared on a reset, panics seen when aborting I/O, and hard errors returned to applications on opens during reset conditions.
 - Changes the evm_versw_undo script to fix no-roll installation and deletion of the EVM version-switched patch.
 - Fixes a problem with the logging of MUNSA reject status messages to the console during boot which could cause a system to boot extremely slowly.
 - Fixes a kernel memory fault from sth_close_fifo() caused by a NULL pointer.
 - Corrects a problem to allow cluster unlinked files to be handled properly during a relocation.
 - Fixes a deadlock problem when deleting devices while the system disk is in error recovery.
 - Fixes POSIX semantics for accessing a "." entry.
 - Closes a race condition between VFS and UFS layer code that causes panic while periodic sync mechanism flushes dirty buffers out to disks.
 - Fixes performance shortcomings in NXM thread replacement.
 - Reduces the number of inputs/outputs to the disk, which reduces the number of audible disk ticks.
 - Corrects a problem where, when NFS mounts using the -o proplist option, disk space is not being freed when files are deleted.
 - Fixes a kernel memory fault in u_seg_global_destroy.
 - Fixes a crash when an AdvFS filesystem reports I/O errors and enters into a domain panic state. AdvFS error cleanup would panic on an invalid pointer and report an "invalid memory read access from kernel mode" panic message.
-

Patch 1367.00 Continued

- Fixes an ISO9660 file system size limitation of 2.1GB and provides full capacity access to DVD-ROM media.
 - Fixes a problem that can prevent certain tape applications from recovering paths to devices that have failed.
 - Prevents a panic from occurring while trying to mount an AdvFS domain. The panic would only appear when the mount command was about to fail.
 - Corrects a problem with reservation conflicts in the TUR recovery loops.
 - Fixes two NFS kernel memory fault panics due to bad NFS server data.
 - Corrects "u_anon_free: page busy panics."
 - Corrects a problem where an I/O can fail back to the application when a HSV110 V2 path failover is performed.
 - Corrects a situation where a tape open with no tape present in the drive can take as long as six minutes to fail.
 - Fixes a VFS namecache race condition where both positive and negative namecache entries can exist.
 - Corrects a problem where low UBC memory conditions cause a hang in AdvFS.
 - Corrects a problem in cluster backups of global root directories or backup of different system disks in a cluster.
 - Fixes an AdvFS alignment fault panic caused by inconsistent AdvFS metadata in a directory. In particular, the directory's entry size is an unaligned value.
 - Creates a condition where, if an I/O fails and it can be helped by an AdvFS-initiated retry, a message is written to the console providing information on how to retry.
 - Fixes an internal value related to the maximum I/O size for a device was being calculated incorrectly.
 - Provides support for the DEGXA Gigabit Ethernet device, including the AlphaServer ES25 onboard 10/100/1000 Ethernet port.
 - Adds the Reserve/Release/Path_Lock feature to tape and changer drivers.
 - Corrects a problem of writing erroneous data to the binary error log file and provides missing header definitions for error interpretation.
 - Provides a variety of domain panic fixes that better capture, explain, and handle domain panics.
 - Fixes a cluster-as-NFS-server chown() problem.
 - Fixes a problem in which the system panics while running applications performing an open of a RAID device and the faulting routine was control_port_open.
 - Fixes a problem where cluster file system I/O and AdvFS domain access causes processes to hang.
 - Fixes a situation in which a system running CD record software experiences I/O timeout errors when writing CDs.
 - Helps avoid a silent infinite loop in vdump by correcting the AdvFS system call OP_GET_BKUP_XTNT_MAP. The call will now return the valid xtntCnt when it fails due to E_NOT_ENOUGH_XTNTS.
 - Fixes a problem where, in certain cases, large files (greater than 30 Gb) suffered extreme fragmentation.
 - Fixes a kernel panic with "bfs_alloc: kernel memory fault."
 - Corrects a problem which had resulted in broadcast or multicast packets being processed multiple times on behalf of a NetRAIN device, once for each backup interface.
 - Fixes a problem that caused the 4.3BSD socket interface to return incorrect values for IOCTL calls accessing IP alias address information.
-

Patch 1367.00 Continued

- Provides support for a cluster patch that corrects a performance issue seen when multiple threads/processes simultaneously access the same file on an SMP (>1 CPU) system.
 - Corrects a problem where, when entering the `hwmgr -view devices` command on a member in a cluster, the device name would not be updated and would be listed as unknown.
 - Corrects a situation where only the `a` and `c` device special files are created when adding a CD-ROM or floppy after boot. Users no longer have to enter `dsfmgr -K` or reboot in order to make all the device special files.
 - Fixes heap and stack limitations in the older operating system versions required for SAP.
 - Prevents several possible system panics and an AdvFS deadlock.
 - Fixes a problem that allowed an application with superuser privileges to cause a system panic when attempting to delete a non-existent connection; for example, when the `trcpkill` program runs while stopping ASU.
 - Fixes an AdvFS AIO read timing issue when reading a fragged file via `directIO`.
 - Prevents a panic when more metadata file space is needed and the disk write to allocate it fails.
 - Removes a restriction where dynamic VMEbus device drivers could only probe one controller per driver. Multiple controllers per driver now configure successfully.
 - Fixes kernel memory faults caused by `ufs_sync_int` accessing an inactivated or de-allocated vnode. This change also corrects a problem with negative block number detection in `ufs_strategy`.
 - Provides support for the pseudo device `/dev/poll` to the kernel, which allows for efficient polling of a large number of file descriptors.
 - Corrects a problem that caused the RFC 2001 Fast Retransmit Algorithm within the kernel to work incorrectly.
 - Makes several changes to the `ee` driver for DE60x Ethernet cards, including the following (these problems affect all Tru64 systems containing DE60x network interfaces):
 - Fixes a race condition that can cause a panic when a transmit timeout occurs.
 - Improves error checking when allocating buffers.
 - Fixes DMA resource allocation to prevent a panic when a machine runs low on DMA resources.
 - Adds a new `ee` subsystem attribute `link_check_interval` that allows the link state polling interval to be tuned for faster failover times when using DE60x interfaces for Link Aggregation.
 - Fixes a `dmapi` problem where `showfile` can show that `dmapi` regions exist when they do not.
 - Provides support for a cluster-specific patch that fixes a race between cluster mounts and file system lookups, and fixes a situation in which file system failover deadlocks.
 - Fixes an NFS readahead performance problem where performance is degraded when reading past 2 GB in a file.
 - Fixes a problem in which `advscan` incorrectly processes concatenated options (for example, `ar` rather than `-a -r`). For instance, if `-ar` is specified, the `-r` option will not be processed.
 - Prevents a lock hierarchy violation from occurring when AdvFS tries to extent a file on a system that is out of memory.
 - Addresses the problem of applications hanging with outstanding I/O during high volume I/O in a cluster environment.
 - Prevents some AdvFS domain panics due to inadequate error handling between the HSG80 and the Tru64 UNIX disk driver.
-

Patch 1367.00 Continued

- Re-enables mountd to support exports file with multiline entry using leading spaces as a continued line indicator. The problem was introduced with a patch that increases support of NFS file mounting from 254 to 1024 entries.
- Corrects a problem with arp messages not being sent on interface static routes.
- Provides the PCI indictment for storage component location to diagnose a PCI adapter failure.
- Resolves a deadlock problem as well as a potential problem with incorrect or inconsistent cluster devts that could occur in a cluster when removing or replacing a device.
- Corrects a problem in which moving the power supply from one slot to another can cause a panic.
- Corrects a possible panic when auditing execve with exec_argp/exec_envp enabled.
- Allows the device special file instance numbers to be reduced to the their lowest possible value and avoid runaway device names.
- Corrects a problem where, under certain conditions, invalidating a portion of a very large file can make the file system appear to be hung. Any program trying to access the file system (for example, issuing the ls command), will hang until the file is invalidated. This will only happen when rt_preempt_opt=1.
- Allows multiple applications utilizing RAID services to send maintenance commands without interfering with each other.
- Prevents different threads on multiple RADs from creating multiple references to the same level 3 page table.
- Corrects a problem involving New_wire_method (light weight wiring), sometimes referred to as the Oracle connect problem or Oracle performance problem.
- Prevents the ARMTech kernel malloc invalid size panic.
- Prevents crash dumps on certain systems that use granularity hinted regions.
- Increases the number of file systems that can be mounted from 256 to 1024.
- Fixes audit to generate exportfs_create audit records correctly.
- Prevents a situation in which the disk drives on a DS25 overheat if the system door is removed for too long.
- Prevents a Kernel Memory Fault panic in irefresh while walking the mounted vnode list.
- Fixes a problem resulting in a system panic for applications that directly call nxm_get_bindings.
- Allows users other than root to mount CD-ROM media on directories that they own.
- Prevents the loss of data in files opened for direct I/O when writing in increments smaller than 8k.
- Fixes a problem in network Link Aggregation (LAG) where the MAC address of a LAG interface would change if the link from which it had derived its MAC address went down.
- Fixes a kernel panic with get_xm_page_range_info:kernel memory fault.
- Corrects several issues with page faults and stack object growth.
- Corrects a problem where df was showing negative values for large NFS file systems.
- Corrects a problem in which multi-threaded processes may hang in timed condition waits (pthread_cond_timedwait()) when running realtime system contention scope threads.
- Enables a larger maximum (1,073,741,824) for the inode_hash_size attribute in the UFS subsystem.
- Fixes a problem that was causing the tcp_rad_fasttimo timer to consume excessive amounts of CPU time.

Patch 1367.00 Continued

- Corrects a problem in which a domain ID was not always updated when mounting an AdvFS file system with the -o dual option in a cluster. If the mount -o dual happened on a node other than the node that was serving the original domain, AdvFS did not detect that the domain ID was already active and failed to update the ID for the new domain. The fix is to always create a new domain ID when the -o dual option is used.
 - Corrects a problem introduced in a prior patch that can result in a system panic when outputting through the packet filter.
 - Corrects a problem in which cluster members panics occurred when previously failed paths to a device (for example, hsz80) are restored.
 - Fixes a bug that prevented AdvFS from working correctly with LSM volumes between 1Tb and 2Tb.
 - Causes mkfdmn and addvol to issue a warning if an attempt is made to use an LSM volume greater than 2Tb in size.
 - Fixes a problem in which AdvFS domains from systems running pre-5.0 versions of Tru64 UNIX and mounted on systems running Patch Kit 3 for Version 5.0A would give "corrupted directory entry size" error messages when some files were accessed.
 - Corrects a problem in which systems configured with VX1 graphics card will not return to the console when the halt button is pressed, thereby making the console unusable.
 - Fixes the problem on an I/O slow start after a host or HSV reboot.
 - Fixes an HSV110 snapshot failure problem due to reservation conflict error when multiple paths are present.
 - Fixes a problem that occurs with an AdvFs file system when ACLs are enabled and there is a Default Access ACL on a directory. The permissions of symbolic links created in that directory appear to be incorrect, even though access is unaffected.
 - Fixes a race during AdvFS volume removal that can cause a panic in the bs_osf_complete() routine.
 - Fixes a problem seen with TAHI IPv6 conformance Test #4 for the IPv6 Specification.
 - Removes pre-emption latencies for ICS threads.
 - Fixes a problem with audit data not being displayed by the audit tool.
 - Fixes problems with file object selection/deselection and directories.
 - Fixes NUMA performance issues associated with auditing.
 - Keeps USB from initializing on systems where USB is not supported.
 - Fixes mmap denying a request at address 0 with MAP_FIXED.
 - Fixes (n)madvise with MADV_DONTNEED's handling of shared memory, which currently leads to mismatched memory.
 - Provides the PCI indictment for storage component location to diagnose a PCI adapter failure.
 - Improves msync performance on files that are mapped with the MAP_PRIVATE flag.
 - Improves the process exit procedure for processes that have had the nice command used on them.
 - Fixes the problem of mount/umount failures and panics in MFS, UFS, and FDFS.
 - Eliminates a condition that causes a hang that occurs while unmounting an AdvFS fileset.
 - Corrects a problem of TS202c system installation failures.
 - Corrects a kernel memory fault that can happen when running applications that use the Cray Intra-Node Shared Memory library.
 - Prevents the loss of single system image for an NFS file system mounted from a cluster when there are certain problems communicating with the external NFS server.
-

Patch 1367.00 Continued

- Fixes a panic that can occur when appending to a file when using UFS.
 - Resolves a kernel memory faults in the TCP/IP subsystem.
 - Fixes 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 when the kernel incorrectly closes a socket that causes Sybase 1613 errors to be produced.
 - Enables SmartArray 5300 controller hardware events to be logged to the binary.errlog file during boot. This is useful in diagnosing logical volume state change and physical drive hotswaps that can occur while the system is not booted.
 - Fixes a problem in which fuser is unable to report on all referenced resources. This is a problem when attempting to identify reasons for unmount failures.
 - Fixes a problem with hung NFS threads due to orphaned mbufs.
 - Corrects a problem with the handling of lock acquisition and release ordering when erasing device mapping contained within SCSI_DIDs.
 - Fixes a problem with printing long double values.
 - Fixes a potential memory discrepancy (and subsequent kernel memory fault) by preventing an "mcs_unlock: lock not found" panic when using process-shared with mcs locks.
 - Fixes several problems in the CAM and I/O space.
 - Adds code to print greater than 61 UNIX domain sockets and change file read errors from /dev/kmem to ignore and continue in a running system.
 - Alleviates a temporary hang/pause condition seen when forking or running down an application with several child processes, from a parent process having an extremely large number of unique or discontiguous memory allocations.
 - Corrects a problem in which multithreaded applications may see corruption of the quadword immediately following a buffer that is written by strncpy(), where an update to the quadword by a second thread is lost.
 - Resolves kernel memory faults in the TCP/IP subsystem.
 - 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.
 - Corrects a problem found wherein the rmtmpfiles script would produce errors at startup of the form:

dirclean: lstat failure for starting directory: /.osonly_tmp/: No such file or directory
 - Fixes a problem with the Smart Array driver that could cause a system hang to occur during error recovery when I/O is active.
 - Modifies AdvFS usage of thread_block to ensure that calls to thread_block will not hang the system.
 - Fixes several IPMI-related problems, including the following:
 - Broadcast "Get Device ID" IPMI command
 - OS power down support (shutdown -p)
 - Erroneous fields in 686 OS-detected environmental machine check logout frame
 - Unusually large number of 686 sensor timeouts with heavy system load
 - IPMI always reports -48v sensors as broken, seen as "redundant power supply failed" messages
 - IPMI memory leak
-

Patch 1367.00 Continued

- Fixes an AdvFS asynchronous direct I/O problem that could cause a thread to hang.
 - Fixes a rmvol E_PAGE_NOT_MAPPED error.
 - Eliminates an ENO_MORE_BLKES error seen when COW'ing to a clone file while an rmvol is in progress.
 - Corrects a problem in AdvFS where it avoids a potential stranded log record in memory that does not get out to disk by fixing a race condition.
 - Prevents a device (for example a disk) from getting two sets of device special files. This problem can occur in a cluster if hundreds of new devices are being found on multiple systems while these systems are booting at the same time.
 - Corrects a problem with a hwmgr delete while a SCSI scan is in progress.
 - Corrects a problem in which compression is not turned off when a device is accessed through the noncompression device special file (/dev/tape/tapex) after having accessed the same device through the compression device special file (/dev/tape/tapexc).
 - Fixes segmentation errors that can occur when running SAS.
 - Fixes a panic caused by a problem within the swapping subsystem.
 - Allows the auditing of login and su events based in part on what is in 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.
 - Fixes a problem in which camreport may report negative device IDs.
 - Fixes a multithread timing window in malloc and free where the list of free chunks could become corrupted, resulting in a segfault.
 - Prevents the hardware management cluster database from being reset.
 - Fixes a problem encountered where a truncated AdvFS file erroneously zeroed data for the remaining leading segment of the file.
 - Corrects a kernel memory fault panic that occurs when running an application that uses the Cray Intra-Node Shared Memory library (Sierra systems only).
 - Corrects a problem that could result in the panic of a cluster member or inconsistent data when the sbcompress_threshold configurable is set.
 - Corrects a problem where, under indeterminate circumstances, when drivers are unloaded or unconfigured, the system may crash if certain kmem_debug flags are set (particulary 0x40 or 0x01). The crash stack backtrace will show remove_dynamic_struct as the problem routine.
 - 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.
 - Prevents a panic in fifo_write with the panic message "NULL fifo_bufhdr append pointer".
 - Corrects the erroneous reporting of success, when attempting to write beyond the file size limit using synchronized I/O.
 - Corrects the calculation of _PC_FILESIZEBITS, which is used by the operating system for pathconf file characteristics.
 - Fixes a system panic in AdvFS when an I/O error occurs in the property list component. The panic string seen is "msfs_pl_cur_to_pnt(1) - failed to ref page".
 - Fixes a problem in the CAM subsystem where it would print out "bad block number" to the error log on a recovered read error. The string has been changed to "block number."
 - Fixes problems with NUMA disk statistics.
-

Patch 1367.00 Continued

- Fixes various small problems in dsfmgr.
 - Prevents a potential process (not system) hang seen when a system comes under heavy memory load with monolithic memory use (gigabyte-scale single objects).
 - Removes latencies for ICS threads.
 - Introduces type checking of attributes when registering components with the hardware manager.
 - Corrects a problem where, under high DMA resource utilization, dma_map_load() may generate an invalid bus address in a scatter/gather entry if the device is using 64-bit addressing, which could result in a system crash or data loss.
 - Changes CHIM to fix Ignore Wide Residue fix and Kernel Memory Fault panic.
 - Fixes a potential problem with modifying files via directIO when there is a clone fileset.
 - Fixes three problems and adds support for one new feature in the alt driver for DEGPA Gigabit Ethernet adapters. These problems affect all Tru64 systems containing DEGPA network interfaces.
 - Fixes a situation where mounting a valid CD-ROM the first time fails with the message "no valid file system exists on this partition."
 - Adds an event which indicates that the soft or hard error count has changed on the device indentified in the event.
 - Prevents a cluster AdvFS panic after a disk array controller restart.
 - Fixes a problem where a user wire request (mlock) fails on NUMA machines.
 - Prevents a "u_anon_free: page busy" system panic.
 - Prevents the system panic "PWS_CCB_QUE_REMOVE: ccb not on any list" caused by a device or bus reset occuring during the execution of a command to a media changer device, such as a tape library.
 - Corrects a failure in the safe_open() routine which caused symbolic links given by a relative path from the current working directory sometimes to give ENOENT errors incorrectly.
 - Fixes a rare case of a thread blocking when waiting for memory.
 - Corrects performance issues when accessing a file with direct I/O enabled.
 - Allows the Event Manager daemon, evmd, to stop listening on its default tcp port 619. This capability is not available for clustered systems.
 - Corrects invalid hwmgr show component inconsistency.
 - Fixes a KMF problem that can happen if some nodes in cluster are rebooted and a device is shared by all the nodes.
 - Fixes the counting of files in AdvFS, which was reporting no available inodes when there were plenty.
 - Resolves a problem of not being able to view files on some CDROM media that is created by third party software.
 - Replaces the system panics caused by "Can't clear bit twice" with a domain panic.
 - Fixes a problem when using the getrusage function where the memory usage could be incorrectly reported.
 - Fixes a problem of an occasional deadlock during process exit for multithreaded processes.
-

Patch 1367.00 Continued

- Fixes a memory leak in the NFS server encountered when it receives malformed packets.
 - Fixes a problem in which, when using the hardware manager to show attributes, the LONG_MAX and LONG_MIN values are displayed incorrectly.
 - Fixes a problem in the kernel network subsystem that caused a kernel memory fault panic in the routine m_adj().
 - Prevents the memory troller from starting on platforms with aluminum EV68 CPUs.
 - Fixes potential quota errors during recovery.
 - Provides a configurable setting that will cause an error return for any read of tape from a tape that requests less the full amount of data in the tape block.
 - Fixes the problem of a kernel memory fault panic in the IP multicast loopback code.
 - Eliminates false directory lookup warning messages generated by an incorrect comparison caused by mismatched file ID variable types and slightly improves client caching performance.
 - Fixes a problem in which rebooting immediately after entering a hwmgr -redirect scsi command results in a boot to single user mode with the following error being displayed:
bcheckrc: Device Naming failed boot configure or verify Please correct the problem and continue or reboot INIT: SINGLE-USER MODE
 - Corrects the behavior of the NFS client negative lookup result cache.
 - Corrects problems of audit_tool supplying incorrect or insufficient data about an audit event.
 - Prevents panics caused by bad arguments to system calls.
 - 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.
 - Fixes a problem in the VM subsystem that could cause a crash with the panic string "vm_page_ssm_unwire."
 - Allows for the proper initialization of the member cache in the set descriptor.
 - Helps avoid a potential lock timeout in AdvFS that could cause a panic with the panic string "mcs_lock: time limit exceeded."
 - Corrects mount(), df(), and memory mapping problems that may prevent users from accessing data on some DVD media, or information about the mounted media.
 - Provides more helpful AdvFS informational messages:
 - Advscan now 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 now includes the location of a file that will help you translate the error number into an error message.
 - Prevents false invalid Inquiry data error reports during boot.
 - Fixes mmap accepting negative lengths, which may lead to a "malloc: invalid size" panic.
 - Fixes a problem in which a taso-compiled binary is unable to allocate more memory after performing a series of mmap calls.
 - Corrects the problem where /sbin/ddr_config does not accept values for ReadyTimeSeconds larger than 255. The new limit is 86400 seconds (24 hours).
 - Fixes excessive FIDS_LOCK contention observed when large numbers of files are using system based file locking.
 - Fixes the reporting of device monitoring events and hardware errors during disk recovery from the disk driver to the binary errlog.
-

Patch 1367.00 Continued

- Forces a domain panic instead of a system panic if AdvFS metadata is discovered to be incorrect in `frag_group_dealloc`.
 - Fixes a problem in which offlining a CPU with one or more bound processes can lead to a "malloc_check_checksum: memory pool corruption" panic.
 - Fixes a problem in which `mmap` memory locked with `mlockall()` using the `MCL_FUTURE` flag does not actually become wired automatically.
 - Prevents `addvol` from adding invalid disks into a domain.
 - Fixes an extended regular expression problem where the interval expression "{m,n}" is handled incorrectly.
 - Fixes a problem in the cluster where the non-default alias mounting feature was disabled.
 - Fixes a kernel memory fault panic in AdvFS that could cause a panic from the `imm_page_to_xtnt()` routine.
 - Fixes a problem in which an NFS server exports files on a third-party file system can result in a system crash.
 - Fixes a problem that occurs when a mounting cluster root if the cluster root domain devices are private to different cluster members, causing the cluster to hang when attempting to boot. With this fix, the cluster will boot with a warning to the console. Although this configuration is not recommended, the cluster should be bootable. The problem occurs with non-LSM cluster root domains.
 - Corrects a memory leak in the VM subsystem.
 - Corrects a potential deadlock in hardware configuration subsystem.
 - Fixed the `audit_tool` search algorithm to differentiate between priviled and non-priviled UIDs, and to allow regular expressions in string searches.
 - Fixes a problem when monitoring I/O via `advfsstat`.
 - Installs DECthreads V3.18-150, which is the latest support version of the HP POSIX Threads library for Tru64 UNIX V5.1A. Previous releases of this patch installed the following versions:
 - DECthreads V3.18-148
 - DECthreads V3.18-144
 - DECthreads V3.18-141, which specifically addressed problems with the pre-emption of the symbolic name `table()` by application code, and the alignment of the Stack Pointer in user-created threads.
 - DECthreads V3.18-138, which specifically addressed a problem that could arise when using recursive mutexes with condition variables.
 - DECthreads V3.18-133
 - Adds SCSI reserve/release support to `mt` to assist open SAN tape management.
 - Corrects problems in the `aha_chim` driver that could result in bus hangs, panics, and inappropriate access of freed memory during a high rate of bus resets.
 - Adds enablers for Smart Array controller.
-

Patch 1367.00 Continued

- Fixes a problem of a controller reset being issued on an idle system due to a thread wake-up signal being missed. This problem may occur when using the Smart Array 5300 series controllers.
 - Fixes a potential floating point error in threaded applications.
 - Provides support for DEGXA Gigabit Ethernet, including ES25 onboard 10/100/1000 port and upgrades the driver for systems with existing support.
 - Fixes a race condition introduced by a previous patch.
 - Adds support for binlog text messages when an environmental system event occurs with an AlphaServer DS20L system. In addition, because of the DS20L's thermal sensitivity, the environmental thermal thresholds have been lowered.
 - Updates ddr.mod to support New Hardware Devices V6 (NHD-6) devices.
 - Provides the V1.07 release of the ciss driver, which is the mandatory minimum version to support the Smart Array 5300 Controller.
 - Fixes a flaw in the NFS server that could cause it to crash upon reception of malformed input.
 - Addresses problems with the NFS server in which a crash can occur with a concurrent read and truncate on an AdvFS file or with malformed or malicious READDIR[PLUS] version 3 RPCs.
 - Address a condition in which AdvFS domain panics would occur during HSZ and HSG failovers.
 - Corrects a problem in which some networking applications, especially X.25 and X.29, stopped working after the installation of Patch Kit 3 because of interactions with security-related fixes and their relationship with fstat.
 - Prevents CS Cluster issues with the DCE/DFS file system when pages are being flushed as part of a vnode.
 - Fixes various problems with the bcm driver for DEGXA Gigabit Ethernet that can cause crashes.
 - Provides the V1.08 release of the ciss driver.
-

Number: Patch 1368.00

Abstract: Correct improper file or privilege management

State: Supersedes Patches 181.00, 486.00, 488.00, 191.00, 2.00, 121.00, 241.00, 243.00, 400.00, 401.00, 402.00, 403.00, 404.00, 405.00, 406.00, 407.00, 408.00, 409.00, 410.00, 412.00, 307.00, 302.00, 162.00, 253.00, 255.00, 90.00, 218.00, 303.00, 305.00, 421.00, 422.00, 423.00, 424.00, 426.00, 681.00, 82.00, 683.00, 684.00, 685.00, 686.00, 687.00, 689.00, 807.00, 511.00, 729.00, 541.00, 146.00, 148.00, 126.00, 127.00, 128.00, 129.00, 130.00, 131.00, 132.00, 134.00, 286.00, 6.00, 7.00, 8.00, 9.00, 10.00, 11.00, 12.00, 13.00, 14.00, 15.00, 16.00, 17.00, 18.00, 19.00, 20.00, 21.00, 22.00, 23.00, 24.00, 25.00, 26.00, 27.00, 28.00, 29.00, 30.00, 31.00, 32.00, 33.00, 34.00, 35.00, 36.00, 37.00, 38.00, 39.00, 40.00, 41.00, 42.00, 43.00, 44.00, 45.00, 46.00, 47.00, 48.00, 49.00, 50.00, 51.00, 52.00, 53.00, 54.00, 55.00, 56.00, 57.00, 58.00, 59.00, 60.00, 61.00, 102.00, 63.00, 104.00, 214.00, 236.00, 246.00, 247.00, 248.00, 250.00, 270.00, 271.00, 272.00, 273.00, 274.00, 275.00, 276.00, 277.00, 279.00, 296.00, 298.00, 321.00, 323.00, 325.00, 290.00, 152.00, 93.00, 95.00, 328.00, 329.00, 330.00, 331.00, 332.00, 333.00, 334.00, 335.00, 336.00, 337.00, 338.00, 339.00, 340.00, 341.00, 342.00, 343.00, 344.00, 345.00, 346.00, 347.00, 348.00, 349.00, 350.00, 351.00, 352.00, (353.00, 354.00, 355.00, 356.00, 357.00, 358.00, 359.00, 360.00, 361.00, 362.00, 363.00, 364.00, 365.00, 366.00, 367.00, 368.00, 369.00, 370.00, 371.00, 372.00, 373.00, 374.00, 375.00, 376.00, 377.00, 378.00, 379.00, 380.00, 381.00, 382.00, 383.00, 384.00, 385.00, 386.00, 387.00, 388.00, 389.00, 390.00, 391.00, 392.00, 393.00, 394.00, 395.00, 396.00, 397.00, 399.00, 543.00, 590.00, 592.00, 596.00, 203.00, 594.00, 597.00, 598.00, 599.00, 600.00, 601.00, 602.00, 603.00, 604.00, 605.00, 606.00, 607.00, 608.00, 609.00, 610.00, 611.00, 612.00, 613.00, 614.00, 615.00, 616.00, 617.00, 618.00, 619.00, 620.00, 621.00, 622.00, 623.00, 624.00, 625.00, 626.00, 627.00, 628.00, 629.00, 630.00, 631.00, 632.00, 633.00, 634.00, 635.00, 636.00, 637.00, 638.00, 639.00, 640.00, 641.00, 642.00, 643.00, 644.00, 645.00, 646.00, 647.00, 648.00, 649.00, 650.00, 651.00, 652.00, 653.00, 654.00, 655.00, 656.00, 657.00, 658.00, 659.00, 660.00, 661.00, 662.00, 663.00, 664.00, 665.00, 666.00, 667.00, 668.00, 669.00, 671.00, 672.00, 673.00, 674.00, 675.00, 676.00, 677.00, 678.00, 680.00, 834.00, 835.00, 837.00, 842.00, 846.00, 853.00, 854.00, 855.00, 856.00, 857.00, 858.00, 859.00, 860.00, 861.00, 862.00, 863.00, 864.00, 865.00, 866.00, 867.00, 868.00, 869.00, 870.00, 871.00, 872.00, 873.00, 874.00, 875.00, 876.00, 877.00, 878.00, 879.00, 880.00, 881.00, 882.00, 883.00, 884.00, 885.00, 886.00, 887.00, 888.00, 889.00, 890.00, 891.00, 892.00, 893.00, 894.00, 895.00, 896.00, 897.00, 898.00, 899.00, 900.00, 901.00, 902.00, 903.00, 904.00, 905.00, 906.00, 907.00, 908.00, 909.00, 910.00, 911.00, 912.00, 913.00, 914.00, 915.00, 916.00, 917.00, 918.00, 919.00, 920.00, 921.00, 922.00, 923.00, 924.00, 925.00, 926.00, 927.00, 928.00, 929.00, 930.00, 931.00, 932.00, 933.00, 934.00, 935.00, 936.00, 937.00, 938.00, 939.00, 940.00, 941.00, 942.00, 943.00, 944.00, 945.00, 946.00, 947.00, 948.00, 949.00, 950.00, 951.00, 952.00, 953.00, 954.00, 955.00, 956.00, 957.00, 958.00, 959.00, 960.00, 961.00, 962.00, 963.00, 964.00, 965.00, 966.00, 967.00, 968.00, 969.00, 970.00, 971.00, 972.00, 973.00, 974.00, 975.00, 976.00, 977.00, 978.00, 979.00, 980.00, 981.00, 982.00, 983.00, 984.00, 985.00, 986.00, 987.00, 988.00, 80.00, 418.00, 780.00, 782.00, 1215.00, 990.00, 1217.00, 1248.00, 1320.00, 1322.00, 1323.00, 1324.00, 1326.00, 1338.00, 1339.00, 1340.00, 1341.00, 1343.00, 1355.00, 1360.00, 1362.00, 1363.00, 1364.00, 1365.00

This patch:

- 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. We have corrected this potential vulnerability.
- Fixes a rmvol failure that would be seen as an `E_PAGE_NOT_MAPPED` error when no more space is available for user data migration to another volume in the domain.
- Fixes a potential problem with `vdf` and `showfdmn`, where they could incorrectly display the message:

`showfdmn: No such file or directory`

- Modifies `rmvol` so that error messages reflect why `rmvol` fails.
- Modifies `showfdmn` so that `showfdmn` will not print "Succeeded" on a failure. For example:

`showfdmn: unable to get info for domain 'domain_used'`
`showfdmn: Successful`

Patch 1368.00 Continued

- Advscan incorrectly processes concatenated options (for example, -ar vs. -a -r). For instance, if -ar is specified, the (-r) option will not be processed.
 - Fixes a bug that prevented AdvFS from working correctly with LSM volumes between 1Tb and 2Tb. Also, mkfdmn and addvol will now issue a warning if an attempt is made to use an LSM volume greater than 2Tb in size.
 - Fives more helpful informational messages:
 - Advscan will now say 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 now includes the location of a file that will help you translate the error number into an error message.
 - Fixes many small problems in dsfmgr.
 - Fixes a rmvol E_PAGE_NOT_MAPPED error. Also eliminates an ENO_MORE_BLKs error seen when COWing to a clone file while a rmvol is in progress.
-
-

TruCluster Server Patches

This chapter provides information about the patches included in Patch Kit 4 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 capsule summaries of the purpose of the TruCluster Server patches included in this kit.

Tru64 UNIX patch kits are cumulative. For this kit, this means that the patches and related documentation from patch kits 1 through 3 are included, along with patches that are new to this kit. To aid you in using this document, release notes that are new with this release are listed as (New) in the section head. The beginning of Section 2.2 provides a key for understanding the history of individual patches.

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 successfully install this patch kit:

- A total of ~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.
- Up to ~24 MB of storage space in `/var/adm/patch/backup` may be required for archived original files if you choose to install and revert all patches. See the *Patch Kit Installation Instructions* for more information.
- Up to ~25 MB of storage space in `/var/adm/patch` may be required for original files if you choose to install and revert all patches. See the *Patch Kit Installation Instructions* for more information.
- Up to ~1221 KB of storage space is required in `/var/adm/patch/doc` for patch abstract and README documentation.
- A total of ~184 KB of storage space is needed in `/usr/sbin/dupatch` for the patch management utility.

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

2.1.2 Updates for Rolling Upgrade Procedures

The following sections provide information on rolling upgrade procedures.

2.1.2.1 Unrecoverable Failure Procedure

The procedure to follow if you encounter unrecoverable failures while running `dupatch` during a rolling upgrade has changed. The new procedure calls for you

to run the `clu_upgrade -undo install` command and then set the system baseline. The procedure is explained in the *Patch Kit Installation Instructions* as notes in Section 5.3 and Section 5.6.

2.1.2.2 During Rolling Patch, Do Not Add or Delete OSF, TCR, IOS, or OSH Subsets

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

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

2.1.2.3 Undoing a Rolling Patch

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

2.1.2.6 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.

2.1.2.7 Relocating AutoFS During a Rolling Upgrade on a Cluster

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

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

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

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

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

```
named          application    ONLINE    ONLINE    rye
```

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

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

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

```
# /usr/sbin/caa_stop -f autofs
```
2. Perform the manual relocation of other applications running on that member:

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

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

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

```
# /usr/sbin/caa_startautofs -c target_member
```
2. Continue with the rolling upgrade procedure.

2.1.3 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:

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

2.1.4 Additional Steps Required When Installing Patches Before Cluster Creation

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

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

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

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

```
# /usr/sbin/versw -setnew
```

2. Run `versw` to switch to the new version:

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

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

2.1.5 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`.

2.1.6 Cluster Information for Tru64 UNIX Patch 1367.00

See Section 1.1.12.2 for version switch information related to Tru64 UNIX Patch 1367.00.

2.1.7 Change to gated Restriction — TruCluster Patch 210.00

The following information explains the relaxed Cluster Alias: `gated` restriction, delivered in TruCluster Patch 210.00.

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

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

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

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

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

Note

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

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

- The alias daemon

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

- Multiple interfaces per subnet (for network load balancing)

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

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

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

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

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

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

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

Note

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

- Members run `gated`, but the cluster alias and `aliasd` are independent of it. The administrator has total control over `gated` and its configuration file, `/etc/gated.conf`. This approach is useful for an administrator who wants to enable IP forwarding and configure a member as a full-fledged router.

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

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

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

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

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

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

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

2.1.8 Information for TruCluster Patch 272.00

This section provides information for TruCluster Patch 272.00.

2.1.8.1 Enablers for EVM

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

2.1.8.2 Rolling Upgrade Version Switch

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

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

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

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

Note

Use this procedure only under the following conditions:

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

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

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

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

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

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

2.1.8.3 Restrictions Removed

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

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

the cluster all filesets mounted from that node's boot partition domain will be force unmounted.

2.1.9 CAA and Datastore — TruCluster Patch 242.00

This section provides information about TruCluster Patch 242.00.

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

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

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

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

You are prompted to guide the backward conversion process.

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

2.2 Summary of TruCluster Software Patches

This section provides capsule summaries of the patches in Patch Kit 4 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 3)**
Indicates a patch that was new in the previous Version 5.1A patch kit.
- **Existing**
Indicates a patch that existed in earlier Version 5.1A patch kits.
- **Supersedes Patches ...**
Indicates a patch that was combined (merged) with other patches.

This section provides capsule summaries of the patches in Patch Kit 4 for the TruCluster Server software products.

Number: Patch 27.00

Abstract: Fix for clusterwide wall messages not being received

State: Existing

This patch allows the cluster wall daemon to restart following an EVM daemon failure.

Number: Patch 88.00

Abstract: Fix for cluster hang during boot

State: Supersedes Patch 29.00

This patch addresses a situation where the second node in a cluster hangs upon boot while setting the current time and date with ntpdate.

Number: Patch 121.00

Abstract: Using a cluster as a RIS server causes panic

State: Supersedes Patch 29.00

This patch:

- Fixes a problem that causes a panic when using a cluster as a RIS server.
 - Provides a fix to RIS/DMS serving in a cluster.
-

Number: 136.00

Abstract: Enhancement for clu_autofs shutdown script

State: Existing

This patch makes the /sbin/init.d/clu_autofs script more robust.

Number: 181.00

Abstract: Fixes problems in the DLM subsystem

State: Supersedes patches 39.00, 131.00, 178.00, 179.00

This patch:

- Fixes a panic in DLM when another node in the cluster is halted.
 - Fixes a panic in the DLM deadlock detection code.
 - Fixes a problem where a process using the Distributed Lock Manager can take up to ten minutes to exit.
 - Fixes several DLM related crashes and performance issues.
 - Corrects a problem causing a cluster member panic.
 - DLM was not always returning the resource block information for the sublock even if the sublock was held.
-

Number: Patch 188.00

Abstract: Fixes cluster kernel problem that causes a hang

State: Supersedes patches 70.00 and 186.00

This patch:

- Fixes a panic in the kernel group services when another node is booted into the cluster.
 - Fixes a problem in the cluster kernel that causes the cluster to hang when a member is rebooted into the cluster.
 - Fixes a problem in the cluster kernel that causes one or more members to panic during a cluster shutdown.
-

Number: Patch 195.00

Abstract: Memory Channel API problem causes system hang

State: Existing (Patch Kit 3)

This patch fixes a problem in the Memory Channel API that can cause a system to hang.

Number: Patch 206.00

Abstract: Fixes kernel memory fault in `rm_get_lock_master`

State: Supersedes Patches 11.00, 62.00, 97.00, 145.00, 146.00, 148.00, 203.00, 204.00

This patch:

- Fixes a situation in which one or several cluster members would panic if a Memory Channel cable was removed or faulty.
- Fixes a problem that causes a clusterwide panic with the Memory Channel power off in a LAN interconnect cluster.
- Allows a user to kill a LAN interconnect cluster via Memory Channel.
- Supports Memory Channel usage in a LAN cluster.
- Fixes a problem where the master failover node goes offline during a failover and failing over due to parity errors increasing beyond the limit.
- Fixes a problem in which a bad Memory Channel cable causes a cluster member to panic with a panic string of "rm_eh_init" or "rm_eh_init_prail."
- Provides changes that should make Memory Channel failovers work better and to handle bad optical cables.
- Fixes a problem in which a node booting into a cluster hangs during Memory Channel initialization.
- Fixes a kernel memory fault in `rm_get_lock_master`.
- Fixes a regression for single physical rail Memory Channel configurations.
- Provides a fix to clean up stale data left on an offline physical rail by the Memory Channel driver.

Number: Patch 210.00

Abstract: `aliasd` now interprets NIFF parameters correctly

State: Supersedes Patches 6.00, 7.00, 9.00, 207.00, 208.00

This patch:

- Fixes a problem in which a cluster member loses connectivity with clients on remote subnets.
 - Fixes a problem with `aliasd` not handling multiple virtual aliases in a subnet and IP aliases.
 - Allows cluster members to route for an alias without joining.
 - Fixes a problem with `aliasd` writing illegal configurations into `gated.conf.memembrX`.
 - Fixes a problem with a default route not being restored after network connectivity issues.
 - Fixes a race condition between `aliasd` and `gated`.
 - Fixes a problem with a hang caused by an incorrect `/etc/hosts` entry.
 - Fixes `aliasd_niff` to allow EVM restart.
 - Provides enablers for Compaq Database Utility.
 - Allows `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.
 - Fixes a problem in which with multiple members booting simultaneously `aliasd` would become deadlocked when trying to select the proxy ARP master for cluster aliases. As a result, some aliases could become unreachable because there would be no proxy ARP master.
 - Fixes a problem in which the `aliasd` daemon message "NIFF parameters for interface are too lax" was erroneously output due to the conversion of internal NIFF parameters from seconds to milliseconds. The `aliasd` daemon now interprets NIFF parameters correctly.
-

Number: Patch 212.00

Abstract: Corrects performance issues on starting cluster LSM

State: Supersedes Patch 150.00

This patch:

- Eliminates spurious duplicate error messages when cluster root is under LSM control.
- Corrects performance issues on starting of Cluster Logical Storage Manager with large configurations.

Patch: Patch 242.00

Abstract: Fix for Oracle failure during start-up

State: Supersedes Patches 1.00, 2.00, 3.00, 5.00, 53.00, 54.00, 55.00, 56.00, 57.00, 58.00, 60.00, 66.00, 71.00, 72.00, 74.00, 84.00, 93.00, 95.00

This patch:

- Increases parallelism in CAA event handling.
 - Fixes a problem with CAA in which after the first resource is started CAA cannot start or stop resources, the resource moves to the unknown state, and a core file is left behind by the action of starting and stopping resources.
 - Provides enablers for Compaq Database Utility.
 - Corrects a problem in which datastore may get corrupted due to improper datastore locking. This may occur when multiple CAA CLI commands are run in the background.
 - Corrects a problem in which the `caa_profile` command may complain of failure to create and log EVM events.
 - Corrects a problem in which the `caa_profile -create` command inserts extra attributes, such as REBALANCE, into the profile when used to create an application profile. This will cause CAA GUI to fail to validate the profile.
 - Corrects a problem where the `caa_stat` command can crash, leaving a core file when it receives a SIGPIPE signal. The problem has been known to occur when `caa_stat` output is piped to a command such as `head`.
 - Fixes a problem that occurs when long resource or attribute names are used and the space is not reclaimed correctly when the resource is unregistered.
 - Fixes a caad memory leak caused by `caa_stat -f`.
 - Corrects a problem in which CAA fails to close a TDF after processing a corresponding resource profile. Over time this will lead to reaching the process limit for open file descriptors and will prevent CAA from functioning properly.
 - Changes the `clu_mibs` agent to cause it to retry the connection with the Event Manager daemon (`evmd`) indefinitely until it succeeds. The `clu_mibs` agent's start and stop control has been moved from `/sbin/init.d/clu_max` script to `/sbin/init.d/snmpd` script.
 - Resolves erroneous behavior of resources with dependencies upon other resources (required resources). This solves several problems with starting, stopping, and relocating a resource with dependencies when the resource's start or stop scripts fail, or when relocating during a shutdown.
 - Causes the old datastore to correctly migrate to the new datastore during the rolling upgrade and corrects the problem where no resource information was preserved.
 - Resolves the issue with the default CAA system services (`dhcp` named `cluster_lockd` `autofs`) not running after the installation of the patch kit. In addition to the default CAA system services any previously registered resource would be lost.
 - Prevents member hangs during boot in unusual circumstances that cause the CAA daemon to crash or exit during initialization.
 - Fixes three CAA problems triggered by heavy CAA activity conditions.
-

Number: Patch 244.00

Abstract: Security (SSRT2265, SSRT2265)

State: Supersedes Patches 48.00, 138.00

This patch:

- Provides a warning when an installed patch includes a version switch, which means the patch cannot be removed using the normal patch removal procedure. The warning allows the user to continue with the switch stage or exit `clu_upgrade`.

In addition to the warning prior to the switch stage, this patch also provides additional user information after the user has decided to perform a patch rolling upgrade and has entered the pathname to a patch kit which contains one or more patches requiring a version switch. The additional user information identifies the patches containing the version switch and provides references to the appropriate user documentation.

- Changes a variable to look at 500 files at a time (instead of the current 700) while making tag files, thereby fixing a problem that can occur during the setup stage of a rolling upgrade during tag file creation.
- Fixes a potential security vulnerability that may result in a denial of service (DoS) on systems running TruCluster Server software.

Patch: Patch 246.00

Abstract: Fixes lsm disks and cluster quorum tools problems

State: Supersedes Patches 41.00, 80.00, 173.00, 175.00

This patch:

- Fixes a cluster installation problem of having an LSM disk and a disk media with the same name. Normally, the installation script would not let you install because it was looking at the disk name, not the disk media name.
 - Allows disks over 10 GB to be used as member or quorum disks.
 - Automates the running of `run versw` to resolve issues with version switched patches and cluster installation.
 - Automatically enables IP filtering for the cluster interconnect on cluster installation and member addition.
 - Allows installation on unlabeled disks.
 - Allows the cluster installation to detect layered product kits in `/var` as well as `/usr/var`.
 - Corrects problems with LSM disks and the cluster quorum tools, specifically when a member having lsm disks local to it is down, the quorum tools fail to update quorum, thereby causing other cluster commands to fail.
-

Number: Patch 252.00

Abstract: Fix for ICS panics

State: Supersedes Patches 37.00, 82.00, 132.00, 134.00, 182.00, 183.00, 185.00, 249.00, 250.00

This patch:

- Closes a timing window that can cause Oracle 9i to hang when a remote node in the cluster goes down.
- Fixes a problem in which panics could occur on process termination and in situations involving multiple memory channel adapters.
- Makes the rdginit daemon program safe to execute multiple times on all cluster interconnect types.
- Resolves a problem resulting in an incorrect error status being returned from RdgInit.
- Makes the following changes to Reliable DataGram (RDG):
 - Changes RDG wiring behavior to match VM's fix to wiring GH chunks.
 - Fixes an RDG problem that can result in user processes hanging in an uninterruptable state.
 - Resolves an RDG panic in the RdgShutdown routine.
 - Fixes a problem in which an RDG kernel thread can starve other timeshare threads on a uniprocessor cluster member. In particular, system services such as networking threads can be affected.
- Resolves a potential kernel memory fault when another node is powered off .
- Resolves a potential user process hang under extreme stress conditions.
- Fixes a kernel thread pre-emption problem that can result in panics due to the starvation of other kernel threads.
- Fixes some misleading send/receive byte count statistics.

Number: Patch 254.00

Abstract: Security

State: Supersedes Patch 52.00

This patch:

- Provides enablers for the Compaq Database Utility.
- Fixes a potential security vulnerability where, under certain circumstances, system integrity may be compromised.

Number: Patch 256.00

Abstract: Fix for cluster hang

State: New

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

Number: Patch 259.00

Abstract: Fixes timing problem in the Connection Manager

State: Supersedes Patches 68.00, 257.00

This patch:

- Fixes a problem where node reboots during a clusterwide shutdown would result in difficult-to-diagnose system panics.
 - Fixes connection manager problems that could result in panics.
 - Fixes a timing problem in the Connection Manager that could cause the panics "CNX MGR: COMMIT_TX: INVALID NODE STATE" or "CNX unaligned access".
-

Number: Patch 263.00

Abstract: Fix for cluster panic

State: Supersedes Patches 44.00, 46.00, 189.00, 190.00, 191.00, 193.00, 260.00, 261.00

This patch:

- Fixes a situation where ICS is unable to make progress because heartbeat checking is blocked or the input thread is stalled. The symptom is a panic of a cluster member with the panic string ICS_UNABLE_TO_MAKE_PROGRESS: HEARTBEAT CHECKING BLOCKED/INPUT THREAD STALLED.
- Fixes the problem of a cluster member failing to rejoin the cluster after Memory Channel failover.
- Addresses a panic that occurs when higher priority threads running on a cluster member block the internode communication service Memory Channel transport (ics_ll_mct) subsystem's input thread from execution.
- Fixes numerous panics and hangs with the way a cluster communicates with its nodes.
- Fixes a problem with hang and panics during boot.
- Fixes a problem that causes a panic with the string "rcnx_status: different node."
- Fixes a boot hang of the string:
ics_mct: Node arrival waiting for out of line node down cleanup to complete
- Fixes a clusterwide hang during extensive of Memory Channel traffic.
- Addresses an assertion caused by a bad user pointer passed to the kernel via sys_call.
- Addresses a panic that occurs while another member was going down.

Number: Patch 265.00

Abstract: Fix for cluster alias manager SUItlet

State:New

This patch fixes the problem in which the cluster alias manager SUItlet falsely interprets any cluster alias with virtual={t|f} configured as a virtual alias regardless of its actual setting.

Number: Patch 269.00

Abstract: A node may panic while under load

State: Supersedes Patches 50.00, 200.00, 267.00

This patch:

- Fixes a situation where a cluster shutdown under load on a cluster using a LAN interconnect takes a very long time.
 - Prevents a panic with duplicate incoming connections on boot.
 - Provides a complete and better error message in event of a misconfigured ICS/TCP adapter.
 - Fixes a condition where a node is not allowed to join the cluster after a panic.
 - Addresses a condition where a node may panic while under load.
 - Address a situation involving discarded UDP datagrams that do not come from the correct port.
-

Number: Patch 272.00

Abstract: Improves responsiveness of EINPROGRESS handling

State: Supersedes Patches 12.00, 13.00, 14.00, 15.00, 16.00, 17.00, 18.00, 19.00, 20.00, 21.00, 22.00, 23.00, 25.00, 76.00, 92.00, 98.00, 99.00, 100.00, 101.00, 102.00, 103.00, 104.00, 105.00, 106.00, 107.00, 108.00, 109.00, 110.00, 111.00, 112.00, 113.00, 114.00, 116.00, 140.00, 142.00, 64.00, 86.00, 117.00, 119.00, 43.00, 151.00, 152.00, 153.00, 154.00, 155.00, 156.00, 157.00, 158.00, 159.00, 160.00, 161.00, 162.00, 163.00, 164.00, 165.00, 166.00, 167.00, 168.00, 169.00, 170.00, 172.00, 30.00, 31.00, 32.00, 33.00, 35.00, 78.00, 90.00, 122.00, 123.00, 124.00, 125.00, 126.00, 127.00, 129.00, 144.00, 196.00, 198.00, 202.00, 213.00, 214.00, 215.00, 216.00, 217.00, 218.00, 219.00, 220.00, 221.00, 222.00, 223.00, 224.00, 225.00, 226.00, 227.00, 228.00, 229.00, 230.00, 231.00, 232.00, 233.00, 234.00, 235.00, 236.00, 237.00, 238.00, 240.00, 270.00

This patch:

- Makes AdvFS fileset quota enforcement work properly on a cluster.
 - Corrects a "cfsdb_assert" panic condition which can occur following the failure of a cluster node.
 - Corrects a problem that can cause cluster members to hang while waiting for the update daemon to flush /var/adm/pacct.
 - Prevents a potential hang that can occur on a CFS failover.
 - Allows POSIX semaphores/msg queues to operate properly on a CFS client.
 - Addresses a potential file corruption problem, which could cause erroneous data to be returned when reading a file at a CFS client node. There is also a small possibility that this problem could result in the CFS panic "AssertFailed: bp->b_dev."
-

Patch 272.00 Continued

- Addressses two potential CFS panic conditions that might occur for a DMAPI/HSM managed file system. The panic strings are:
 - Assert Failed: (t)->cntk_mode <= 2
 - Assert Failed: get_recursion_count(current_threa&CMI_TO_REC_LOCK(mi)) == 1
 - Corrects a problem in which a possible panic that could occur if multiple CFS client nodes leave the cluster while a CFS relocate or unmount is occurring.
 - Fixes a problem where a possible KMF panic occurs when executing the command `cfsmgr -a DEVICES` on a file system with LSM volumes.
 - Corrects a CFS problem that could cause a panic with the panic string of "CFS_INFS full".
 - Fixes a problem where a possible CFS panic might occur when a file is opened in Direct I/O mode at the same time it is being truncated by a separate process.
 - Provides enabler support for Enterprise Volume Manager product.
 - Fixes memory a leak in `cfscall_ioctl()`.
 - Provides support for the `freezefs` utility.
 - Fixes a data inconsistency that can occur when a CFS client reads a file that was recently written to and whose underlying AdvFS extent map contains more than 100 extents.
 - Fixes a panic that would occur during the mount of a clusterized file system on top of a non-clusterized file system.
 - Prevents a kernel memory fault panic during unmount in a Cluster or during a planned relocation.
 - Fixes support for mounting other filesets from the `cluster_root` domain in a cluster.
 - Fixes the assertion failure `ERROR != ECFS_TRYAGAIN`.
 - Fixes a race condition during a cluster mount that results in a transient `ENODEV` seen by a name space lookup.
 - Fixes a problem in which a panic on boot could occur if a mount request is received from another node too early in the boot process.
 - Fixes a problem in which a `PANIC: CFS_ADD_MOUNT() - DATABASE ENTRY PRESENT` panic could occur when a node rejoins the cluster.
 - Fixes a race condition in cluster mount support that results in a transient mount failure and a second race that might result in a kernel memory fault panic during mount.
 - Fixes a cluster problem with hung unmounts (possibly seen as hung node shutdowns).
 - Fixes a problem in which a UBC panic could occur when accessing CFS file systems.
-

Patch 272.00 Continued

- Prevents a possible Kernel Memory Fault panic on racing mount update/unmount/remount operations for the same mount point.
 - Fixes a possible race between node shutdown and unmount.
 - Prevents a possible Kernel Memory Fault panic on the mount update on a Memory File System (MFS) and other possible panics when bad arguments are passed to the mount library interface.
 - Prevents the panic "Assert failed: vp->v_numoutput > 0" or a system hang when a file system becomes full and direct asynchronous I/O via CFS is used. A vnode will exist that has v_numoutput with a greater than 0 value and the thread is hung in vflushbuf_aged().
 - Prevents a possible Kernel Memory Fault in function ckidtokgs.
 - Fixes a potential CFS deadlock condition.
 - Corrects the problem of the cfsmgr error "Not enough space" when attempting to relocate a file system with a large amount of disks.
 - Fixes a problem in which possible CFS client node file read failures could occur if the domain storage devices were closed during a previous failure to perform a failover mount on the client node,
 - Fixes support for mounting other filesets from a cluster node's boot partition domain.
 - Addresses a cluster problem that can arise in the case where a cluster is serving as an NFS server. The problem can result in stale data being cached at the nodes which are servicing NFS requests.
 - Fixes a CFS panic that might occur for a DMAP/HSM managed fs: (panic): cfstok_hold_tok(): held token table overflow
 - Fixes a panic "cmn_err: CE_PANIC: ics_unable_to_make_progress: netisrs stalled" in clua.mod due to wait for malloc when memory is exhausted.
 - Fixes a panic in clua_cnx_unregister where a TP structure could not be allocated for a new TCP connection.
 - Fixes problems with cluster alias selection priority when adding a member to an alias.
 - Fixes a problem when the cluster alias subsystem does not send a reply to a client that pings a cluster alias address with a packet size of less than 28 bytes.
 - Allows the cfsstat -i command to execute properly.
 - Fixes a potential Cluster File System deadlock that can occur during CFS failover processing following the failure of a CFS server node.
 - Prevents process hangs on clusters mounting NFS file systems and accessing plock-ed files on the NFS server.
 - Fixes a possible timing window whereby a booting node may panic due to memory corruption if another node dies.
 - Fixes a small window that can cause a clusterwide panic on node reboot in a quorum loss situation.
-

Patch 272.00 Continued

- Fixes a problem in which a cluster member may panic with the panic string "kernel memory fault".
 - Fixes a possible boot hang that could occur if the cluster_root domain consists of LSM volumes whereby the underlying physical storage is nonshared.
 - Prevents a memory leak from occurring when using small, unaligned Direct I/O access (that is, not aligned on a 512 boundary and does not cross a 512 byte boundary).
 - Prevents the cfsmgr command from displaying an erroneous server name when a request is made for statistics for an unmounted file system.
 - Fixes support for Synchronized I/O in clusters.
 - Eliminates erroneous EIO errors that could occur if a client node becomes a server during a rename/unlink/rmdir system call.
 - Corrects a CFS problem that could result in degraded performance when reading at file offsets past 2GB.
 - Corrects a cluster file locking problem that can arise when file systems are exported from the cluster to NFS client nodes.
 - Fixes a CFS problem where file access rights may not appear consistent clusterwide.
 - Fixes a race between cluster mounts and file system lookups.
 - Fixes a problem in which file system failover deadlocks.
 - Corrects a Cluster File System (CFS) performance issue seen when multiple threads/processes simultaneously access the same file on an SMP (>1 cpu) system.
 - Addresses a potential clusterwide hang which can occur in the Cluster File System.
 - Fixes a problem in which file permissions inherited from the default ACL may be different than expected under the following conditions:
 - ACLs are enabled on the system
 - There is a default ACL on a directory
 - A request is issued from a CFS client to create a file within that directory
 - Fixes a problem where cluster file system I/O and AdvFS domain access causes processes to hang.
 - Prevents an infinite loop during node shutdown when using server_only file systems.
-

Patch 272.00 Continued

- Fixes a memory fault panic from `clua_cnx_thread`.
 - Fixes a problem in which an application that uses file locking may experience degraded performance.
 - Provides the I/O barrier code that prevents HSG80 controller crashes (firmware issue).
 - Fixes a situation in which a rebooting cluster member would panic shortly after rejoining the cluster if another cluster member was doing remote disk I/O to the rebooting member when it was rebooted.
 - Allows high density tape drives to use the high density compression setting in a cluster environment.
 - Fixes a kernel memory fault panic that can occur within a cluster member during failover while using shared served devices.
 - Fixes the problem of clusterwide hang because a DRD node failover is stuck and unable to bid a new server for served device.
 - Adds DRD Barrier retries to work around HSx firmware problems.
 - Fixes a problem in which CAA applications using tape/changers as required resources will not come ONLINE (as seen by `caa_stat`).
 - Fixes a problem in which the tape changer is only accessible from member that is the DRD server for the changer.
 - Fixes a problem where an open request to a disk in a cluster fails with an illegal errno (≥ 1024).
 - Fixes a problem where an open to a tape drive in a cluster would take six minutes (instead of two) to fail if there were no tape in the drive.
 - Corrects a problem in which a cluster would hang the next time a node was rebooted after a tape device was deleted from the cluster.
 - Fixes a domain panic in a cluster when a file system is mounted on a disk accessed remotely over the cluster interconnect.
 - Fixes the race condition problem when multiple unbarrierable disks failed at the same time.
 - Fixes a kernel memory fault in `drd_open`.
 - Prevents an infinite loop in `drd_open()`.
 - Fixes several Device Request Dispatcher problems.
 - Provides the required mechanism to remove a rolling upgrade issue with CD-ROM and floppy disk device handling.
 - Fixes a problem in which a cluster or a device can get I/O stuck or that a cluster node may panic after a device has been deleted.
 - Fixes a problem of excessive FIDS_LOCK contention that occurs when large number of files are using system-based file locking.
 - Causes the immediate updating of the attributes on a directory when files are removed by a cluster node that is not the file system server.
 - Fixes a hang condition in Device Request Dispatcher (DRD) when accessing a failed disk.
-

Patch 272.00 Continued

- Prevents a simple_lock: time limit exceeded panic or a Assert Failed: brp->br_fs_svr_out panic that can be seen while executing chfsets on a cluster.
 - Fixes problems in the cluster kernel where a cluster member hangs during cluster shutdown or while booting.
 - Fixes a problem in the cluster kernel where a cluster member panics when a tape device is accessed.
 - Fixes a token problem that could cause an unmount to hang.
 - Fixes a condition that causes the panic "CNX MGR: Invalid configuration for cluster seq disk" during simultaneous booting of cluster nodes.
 - Prevents an infinite loop in drd_open().
 - Fixes a problem in which two nodes leaving the cluster within a short time period would cause I/O on some devices to get stuck.
 - Fixes a problem in which a new device would not be properly configured in a cluster if the device was discovered during a boot.
 - Causes the Device Request Dispatcher (DRD) to retry to get disk attributes when EINPROGRESS is returned from the disk driver.
 - Fixes an issue with ICS (Internode Communication Services) on a NUMA-based system in a cluster.
 - Fixes a possible race condition between a SCSI reservation conflict and an I/O drain that could result in a hang.
 - Adds support for multiple opens to tape libraries/media changers.
 - Alleviates a condition in which a cluster member takes an extremely long time to boot when using LSM.
 - Corrects reference-counting errors that may lead to a panic during cluster mount.
 - Relieves pressure on the CMS global DLM lock by allowing AutoFS auto-mounts to back off.
 - Addresses a potential panic in the Cluster File System that can occur when using raw Asynchronous I/O.
 - Addresses a potential panic in the Cluster File System that can occur when using file system quotas.
 - Fixes kernel memory faults associated with passing in invalid parameters to the mount system call.
 - Fixes the problem of a potential hang when multiple nodes are shutting down simultaneously and server-only file systems are mounted.
 - Fixes the problem of a potential system crash when adding a cluster alias.
 - Improves the responsiveness of EINPROGRESS handling during the issuing of I/O barriers. The fix removes a possible infinite loop scenario which could occur due to the deletion of a storage device.
 - Allows AutoFS auto-UNmounts to back off, thereby relieving pressure on the CMS global DLM lock.
 - Adds data validation checking pertaining to cluster messages involving tokens, to assist in problem isolation and diagnosis.
 - Corrects diagnostic code that might result in a panic during kernel boot.
 - Corrects a problem in which a bus reset causes the loss of quorum, resulting in a cluster hang.
-
-

Revised Reference Pages

This appendix provides changes to reference pages due to patches included in this kit.

A.1 envconfig(8) Update

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

`envconfig(8)`

NAME

`envconfig` - Configures the Environmental Monitoring daemon

SYNOPSIS

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

`/usr/sbin/envconfig start | stop`

`/usr/sbin/envconfig -q`

OPTIONS

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

`-c var=value`

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

`ENVMON_CONFIGURED`

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

`ENVMON_GRACE_PERIOD`

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

`ENVMON_HIGH_THRESH`

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

`ENVMON_MONITOR_PERIOD`

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

`ENVMON_USER_SCRIPT`

Specifies the path of a user-defined script that you want the `envmond` daemon to execute when a high threshold level is encountered. The `envmond` daemon continues to check the environment after

the script has executed and proceeds as needed should the high threshold levels persist.

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

ENVMON_SHUTDOWN_SCRIPT

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

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

start | stop

Turns the envmond daemon on or off after system startup.

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

DESCRIPTION

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

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

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

EXAMPLES

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

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

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



```
env_ps_status = 0
env_supported = 1
```

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

- + At the command prompt, enter the following command:
`/usr/sbin/envconfig -c ENVMON_CONFIGURED=1`

- + Use the `rcmgr` command as follows:
`rcmgr set ENVMON_CONFIGURED 1`

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

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

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

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

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

FILES

`/etc/rc.config*`

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

SEE ALSO

Commands: `envmond(8)`

A.2 `sys_check(8)` Update

Patch 807.00 changes the `envconfig` command. See Section A.2 for an updated `sys_check(8)` reference page.

`syscheck(8)`

NAME

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

SYNOPSIS

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

OPTIONS

`-all`

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

`-debug`

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

`-escalate [xx]`

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

in the `escalate.tar` file. Optionally, you can specify a number (xx) with the `-escalate` option to define a crash number.

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

-evm

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

- + 0 - Information only.
- + 100 - Note
- + 200 - Tuning Note
- + 300 - Tuning Suggestion
- + 400 - Operational
- + 500 - Warning

-frame

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

-name name

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

-dir name

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

-help or (-h)

Outputs help information.

-nohtml

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

-noquick

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

-perf

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

-v Displays the `sys_check` version number.

-warn

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

-nowarn

Executes only the data gathering pass.

DESCRIPTION

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

The `sys_check` utility also performs an analysis of operating system parame-

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

When used as a component of fault diagnosis, `sys_check` can reduce system down time by as much as 50% by providing fast access to critical system data. It is recommended that you run a full check at least once a week to maintain the currency of system data. However, note that some options will take a long time to run and can have an impact on system performance. You should therefore choose your options carefully and run them during off-peak hours. At a minimum, perform at least one full run (all data and warnings) as a post-configuration task in order to identify configuration problems and establish a configuration baseline. The following table provides guidelines for balancing data needs with performance impact.

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

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

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

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

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

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

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

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

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

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

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

NOTES

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

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

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

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

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

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

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

If `DECevent` is not present, the `sys_check` utility issues a warning message

as a priority 500 EVM event and attempts to use uerf instead. In future releases, Compaq Analyze will also be supported on certain processors.

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

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

EXIT STATUS

The following exit values are returned:

- 0 Successful completion.
- >0 An error occurred.

LIMITATIONS

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

HSZ controller-specific limitations include the following:

HSZ40 and HSZ50 controllers:

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

HSZ70, HSZ80 and G80 controllers:

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

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

EXAMPLES

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

6. The following command starts the SysMan Menu `config_report` task from the command line:

```
# /usr/sbin/sysman config_report
```

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

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

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

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

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

```
advfs:
AdvfsCacheMaxPercent=8
```

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

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

ENVIRONMENT VARIABLES

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

TMPDIR

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

LOGLINES

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

BIGNUMFILE

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

BIGFILE

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

VARSIZE

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

RECOVERY_DIR

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

ADHOC_DIR

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

TOOLS_DIR

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

FILES

`/usr/sbin/sys_check`

Specifies the command path.

Note

This file may be a symbolic link.

`/usr/sbin/*`

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

Note

These files may be symbolic links.

The `sys_check` utility reads many system files.

SEE ALSO

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

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

Books: *System Administration*, *System Tuning*

A.3 `sys_attrs_netrain(5)`, `niffmt(7)`, `niffconfig(8)`, and `ifconfig(8)` Updates

The installation of Patch 1107.00 results in changes to the `sys_attrs_netrain(5)`, `niffmt(7)`, `niffconfig(8)`, and `ifconfig(8)` reference pages. To review these changes see the following sections:

A.3.1 `sys_attrs_netrain(5)`

`nr_timeout_dead_interface`

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

Minimum value: 0.5 (seconds)

`nr_timeout_o`

Minimum value: 1.1

`nr_timeout_t`

Minimum value: 0.5

You can specify decimal values (for example, 2.5 or 0.8) for `nr_timeout_dead_interface`, `nr_timeout_o`, and `nr_timeout_t`. When you recon-

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

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

SEE ALSO

`sys_attrs(5)`, `nifftmt(7)`

Network Administration: Connections

A.3.2 `nifftmt(7)`

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

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

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

```
t2 >= t1 + 2dt
```

```
t1 >= 0.5
```

```
t2 >= 1.1
```

```
dt >= 0.2
```

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

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

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

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

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

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



```

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

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

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

```

A.3.3 niffconfig(8)

SYNOPSIS

```

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

```

-d dt

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

-o t2

Specifies the total number of traffic-free seconds that must elapse before the traffic monitor thread determines that a network interface has failed. This number must be at least the sum of the t1 and two times dt. That is, given the default time period for dt (5 seconds) and t1 (20 seconds), the t2 value must be at least 30 seconds. The

default time period for t2 is 60 seconds. If t2 is not specified, niffconfig uses the default.

-m Modifies the timing parameters of an interface that is already being monitored. Typically, this option is specified along with one or more of **-t t1**, **-d dt**, or **-o t2** options. If none of these parameters are specified, the default value is used. You cannot specify the **-m** option with the **-a**, **-s**, **-r**, **-u**, or **-v** options.

-t t1
Specifies the time period, in seconds, that the traffic monitor thread delays between reads of the interface counters when the network is running normally. The default time period is 20 seconds. If t1 is not specified, niffconfig uses the default.

-v Displays the status, timer values, and description (verbose mode) of all interfaces currently being monitored to standard out (stdout). See niffmt(7) for a definition of each of the parameters.

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

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

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

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

EXAMPLES

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

A.3.4 ifconfig(8)

The `monitor` section should be removed:

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

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

A.4 lag(7) Update

Patch 1367.00 changes the information in the `lag(7)` reference page.

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

NAME

`lag` - Link aggregation (also called trunking) introductory information

DESCRIPTION

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

Using link aggregation provides the following capabilities:

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

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

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

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

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

Link Aggregation Configuration

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

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

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

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

Fault Tolerance

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

link aggregation group and redistributes traffic among the remaining ports.

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

Load Sharing

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

Destination IP Address (dstip)

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

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

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

Destination MAC address (dstmac)

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

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

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

Transport Port number (port)

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

This algorithm can utilize the combined bandwidth of a link aggregation group in environments where traffic is destined to a single IP or MAC address, but is exchanged between a number of different TCP or UDP port number pairs (for example, a server-to-server interconnect).

Round Robin (roundrobin)

The port is selected on a rotating basis.

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

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

RESTRICTIONS

The following restrictions apply:

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

RELATED INFORMATION

Commands: lagconfig(8)

System Attributes: sys_attrs_lag(5)

Files: inet.local(4)

Technical Overview

Network Administration: Connections

A.5 lagconfig(8) Update

Patch 1367.00 changes the information in the lagconfig(8) reference page.

lagconfig(8) lagconfig(8)

NAME

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

SYNOPSIS

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

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

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

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

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

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

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

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

OPTIONS

-c Creates a link aggregation group virtual interface. You can specify the following attributes to this option. If you specify more than one

attribute, separate them with commas:

lag=interface-id

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

key=value

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

dist={dstip | dstmac | port | roundrobin}

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

macaddr=address

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

-p port

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

lag=interface-id

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

key=value

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

-d port

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

-s lag=interface-id

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

DESCRIPTION

The lagconfig command allows you to perform the following tasks:

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

Link aggregation, or trunking, enables administrators to combine one or

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

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

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

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

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

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

EXAMPLES

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

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

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

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

Note

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

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

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

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

```
ifconfig lag0 10.1.2.3 up
```

DIAGNOSTICS

```
lagconfig: subsystem error: Invalid argument
```

You attempted to add a port (interface) to a link aggregation group and the port is UP. Mark the interface DOWN with the `ifconfig` command and try to add the port again.

SEE ALSO

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

Interfaces: lag(7)

System Attributes: sys_attrs_lag(5)

Network Administration: Connections

A.6 wol(8) Update

The installation of this patch changes the information in the wol(8) reference page.

wol(8)

NAME

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

SYNOPSIS

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

OPTIONS

nw_interface

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

OPERANDS

hw_address

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

DESCRIPTION

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

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

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

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

```
>>> set wol_enable on
>>> init
```

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

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

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

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

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

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


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

Note

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

RESTRICTIONS

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

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

EXIT STATUS

0 (Zero)

Success.

>0 An error occurred.

ERRORS

+ Error detecting default interface

Explanation:

The wol utility cannot automatically detect a default network interface.

User Action:

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

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

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

Explanation:

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

User Action:

Specify the hardware network address correctly.

EXAMPLES

1. The following example shows a simple use of the wol utility, where the host system detects the first configured network interface and prompts for confirmation:
/usr/sbin/wol 00-02-56-00-03-29
No sending device specified, using tu0, continue? (y/n) y
2. The following example shows the same use of the wol utility, where the user declines confirmation of the selected network interface:
/usr/sbin/wol 00-02-56-00-03-29
No sending device specified, using tu0, continue? (y/n) n
Aborting...
3. The following example explicitly specifies a network interface:

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

ENVIRONMENT VARIABLES

wol_enable

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

Note

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

FILES

/usr/sbin/wol

Wake-on-LAN utility.

SEE ALSO

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

New Hardware Delivery Release Notes and Installation Instructions

System Administration