

Tru64 UNIX and TruCluster Server Version 5.1A

Patch Summary and Release Notes for Patch Kit 6

November 2003

This manual describes the contents of Patch Kit 6 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.

© Copyright 2003 Hewlett-Packard Development Company, L.P.

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

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

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

Contents

About This Manual

1 Tru64 UNIX Patches

1.1	Release Notes	1-1
1.1.1	Required Storage Space	1-1
1.1.2	Changes to Reference Pages	1-1
1.1.3	Potential NFS Duplicate Request Cache Scalability Limitation with High Loads and Uncharacteristic File Access Behavior on Clustered NFS Servers (new)	1-2
1.1.4	Tuning the NFS Server Duplicate Request Cache (new)	1-4
1.1.5	Moving Files in DMAPI-Enabled Filesets Across Partitions — Patch 2084.00 (new)	1-4
1.1.6	Performance of hwmgr Commands on Large System Configurations (new)	1-4
1.1.7	Packetfilter Enhancements — Patch 2084.00 (new)	1-5
1.1.8	Changes to vdump and vrestore Allow Larger Record Sizes — Patch 2064.00	1-5
1.1.9	Potential Problem Patching a System with WLS Subsets Installed	1-6
1.1.10	Security Feature — Patch 2084.00	1-6
1.1.11	Changes to tar, pax, and cpio Behavior — Patch 2088.00	1-8
1.1.12	Ignore Special Instruction for Patch 1924.00	1-8
1.1.13	Russian Keyboard — Patch 1197.00	1-8
1.1.14	Problem Seen on Systems with Smart Array Controller	1-8
1.1.15	Problem with lockmode=4 on AlphaServers with QLogic SCSI Disk Controllers	1-9
1.1.16	Updates to sys_check	1-9
1.1.16.1	TMPDIR Variable	1-9
1.1.16.2	sys_check Version 125 Web Kit	1-9
1.1.17	Support for SDLT160 Tape Device	1-10
1.1.18	New libots3 Library — Patches 1146.00 and 1148.00	1-11
1.1.19	New esmd Daemon — Patch 252.00	1-11
1.1.20	Command Updates and Other Changes — Patch 2084.00	1-11
1.1.20.1	Updates to sh, csh, and ksh	1-11
1.1.20.2	sh noclobber Option and > , >> Constructs Added	1-11
1.1.20.3	ksh noclobber Behavior Clarified	1-12
1.1.20.4	csh noclobber Behavior Clarified	1-12
1.1.20.5	Updated mkdir System Call and Command	1-12
1.1.20.6	Enabling the /dev/poll Function	1-12
1.1.20.7	Removal of Version-Switched Patch	1-13
1.1.20.8	New ee Attribute	1-13
1.1.20.9	Support for Network Link Aggregation	1-13
1.2	Summary of Base Operating System Patches	1-13

2 TruCluster Server Patches

2.1	Release Notes	2-1
-----	---------------------	-----

2.1.1	Required Storage Space	2-1
2.1.2	Removing Some Patches Can Cause Problems	2-1
2.1.3	Patch Removal Causes Login Error	2-2
2.1.4	Updates for Rolling Upgrade Procedures	2-2
2.1.4.1	Problem When Undoing Roll with Worldwide Languages Installed	2-2
2.1.4.2	Order for Rolling NHD6 and Patch Kit 6	2-2
2.1.4.3	Unrecoverable Failure Procedure	2-2
2.1.4.4	During Rolling Patch, Do Not Add or Delete OSF, TCR, IOS, or OSH Subsets	2-3
2.1.4.5	Undoing a Rolling Patch	2-3
2.1.4.6	Ignore Message About Missing ladefug.cat File During Rolling Upgrade	2-3
2.1.4.7	clu_upgrade undo of Install Stage Can Result in Incorrect File Permissions	2-3
2.1.4.8	Missing Entry Messages Can Be Ignored During Rolling Patch	2-4
2.1.4.9	Relocating AutoFS During a Rolling Upgrade on a Cluster ..	2-4
2.1.5	When Taking a Cluster Member to Single-User Mode, First Halt the Member	2-5
2.1.6	Additional Steps Required When Installing Patches Before Cluster Creation	2-5
2.1.7	Problems with clu_upgrade switch Stage	2-6
2.1.8	Cluster Information for Tru64 UNIX Patch 2084.00	2-6
2.1.9	Change to gated Restriction — Patch 336.00	2-6
2.1.10	Version Switch Warning Added — Patch 351.00	2-7
2.1.11	Information for Patch 371.00	2-8
2.1.11.1	Enablers for EVM	2-8
2.1.11.2	Rolling Upgrade Version Switch	2-8
2.1.11.3	Restrictions Removed	2-9
2.1.12	CAA and Datastore — Patch 338.00	2-9
2.2	Summary of TruCluster Software Patches	2-9

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), and niffconfig(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.4	wol(8) Update	A-12

About This Manual

This manual contains information specific to Patch Kit 6 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.
- *Tru64 UNIX and TruCluster Server Patch Kit Installation Instructions*
- *Patching Best Practice*
- The `dupatch(8)` reference page, which describes the use of `dupatch` from the command line. This reference page is installed when you install the `dupatch` tools.
- *Tru64 UNIX Installation Guide*
- *Tru64 UNIX System Administration*
- *TruCluster Server Cluster Installation*
- *TruCluster Server Cluster Administration*
- Release-specific installation documentation

Patch Process Resources

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

- To obtain the latest patch kit for your operating system and cluster:
`http://www.itrc.hp.com/service/patch/mainPage.do`
- To view or print the latest version of the *Patch Kit Installation Instructions* or the *Patch Summary and Release Notes* for a specific patch kit:
`http://h30097.www3.hp.com/docs/patch/index.html`

- To visit our main support page:
<http://h71025.www7.hp.com/support/home/index.asp>
- To visit the Tru64 UNIX homepage:
<http://h30097.www3.hp.com/>

Reader's Comments

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

You can send your comments in the following ways:

- Fax: 603-884-0120 Attn: HCTO Information Development, ZK03-3/Y32
- Internet electronic mail:
readers_comment@zk3.dec.com

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

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

Please include the following information along with your comments:

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

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

Tru64 UNIX Patches

This chapter provides information about the patches included in Patch Kit 6 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 5 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 87.9 MB of storage space in /var/adm/patch/backup may be required for archived original files if you choose to install and revert all patches. See the *Patch Kit Installation Instructions* for more information.
- Approximately 90.6 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 2.4 MB of storage space is required in /var/adm/patch/doc for patch abstract and README documentation.
- Approximately 184 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 Changes to Reference Pages

The following table lists reference pages that have been revised since the Version 5.1A release of the operating system and are installed by patches in this kit.

Changes to several reference pages that are not revised via patches are described in Appendix A.

Reference Page	Patch ID or Section
chatr(1)	Patch 1624.00
ckfsec(1)	Patch 824.00
dirclean(8)	Patch 763.00
envconfig(8)	Section A.1
exportfs(2)	Patch 763.00
ifconfig(8)	Patch 1370.00
javaexecutedata(8)	Patch 1701.00
kdbx(8)	Patch 1611.00
lag(7)	Patch 1651.00
lagconfig(8)	Patch 1755.00
mktemp(1)	Patch 701.00
mktemp(3)	Patch 705.00
netstat(1)	Patch 2066.00
niffconfig(8)	Section A.3.3
niffmt(7)	Section A.3.2
poll(7)	Patch 765.00
route(8)	Patch 2066.00
safe_open(3)	Patch 731.00
sys_attrs_ee(5)	Patch 1717.00
sys_attrs_dli(5)	Patch 1669.00
sys_attrs_inet(5)	Patch 2066.00
sys_attrs_net(5)	Patch 2066.00
sys_attrs_netrain(5)	Section A.3.1
sys_attrs_proc(5)	Patch 1701.00
sys_check(8)	Section A.2
volrestore(8)	Patch 2072.00
volsave(8)	Patch 2072.00
wol(8)	Section A.4

1.1.3 Potential NFS Duplicate Request Cache Scalability Limitation with High Loads and Uncharacteristic File Access Behavior on Clustered NFS Servers (new)

Repeated simultaneous overwriting of many files can cause retransmitted writes to be processed after recent writes of a file to the same location. This problem occurs more often on systems configured with a LAN cluster interconnect than on those configured with Memory Channel.

This behavior is inherent in the "stateless" design of NFS. Although the behavior has been mitigated via a "duplicate request cache" that replays old replies instead of reexecuting retransmitted requests, extremely heavy loads on large systems can overwhelm the cache when requests are stalled. Customers are unlikely to see problems because applications rarely rewrite files almost immediately.

If the problem occurs, the NFS server displays the following message several times a minute on the system console, indicating that the NFS server is being overwhelmed with requests :

```
"NFS server xxx not responding"
```

When an "overwhelmed duplicate request cache" condition has occurred, the NFS client will display multiple occurrences of either of the following messages:

```
NFS3 server xxx not responding still trying
NFS3 server xxx ok
```

```
NFS2 server xxx not responding still trying
NFS2 server xxx ok
```

This indicates that the client is observing transient unresponsive periods at the server. This is the only notification that the client will display if the server's duplicate request cache becomes overwhelmed. When the client detects this behavior, it increases the retransmission interval until it gets a response from the server. This behavior is generally undistinguishable from the server going up and down, except that the messages are displayed with such frequency that the server system/member cannot have gone down and then come back up in that short an interval.

You can minimize the likelihood of these problems as follows:

- Avoid congestion on your LAN and cluster interconnect.
- Ensure your servers have enough excess capacity to respond quickly to NFS requests that modify the file system (writes, file and directory creation, and so forth.)
- Increase the size of the server's duplicate request cache when the `nfsstat` command shows a large number of retransmits to clients. For instructions on increasing the size of the cache, see Section 1.1.4.

You can monitor the number of NFS retransmissions using the `nfsstat -c` command. . The `retrans` field indicates the number of retransmissions. A retransmission rate higher than 2% indicates a potential problem.

The following example shows the output from the `nfsstat -c` command. The retransmission fields are marked with asterisks (*). This example is of a client workstation in a typical environment.

```
% nfsstat -c
```

```
Client rpc:
```

```
tcp:  calls  badxids  badverfs  timeouts  newcreds
      0      0      0      0      0
      creates  connects  badconns  inputs  avails  interrupts
      0      0      0      0      0      0
udp:  calls  badxids  badverfs  timeouts  newcreds  *retrans*
      224518870  959      0      101985      0      0
      badcalls  timers  waits
      102013  110894      0
```

```
Client nfs:
```

```
calls  *retrans*  badcalls  nclget  nclsleep  ndestroys  ncleans
224414222  4248      28  224414282      0      6219      224408063 \
```

If an overwhelmed duplicate request cache condition occurs, we recommend you perform one or more of the following tasks:

- Ensure that there are short periods of idle time on the I/O subsystem and network links.
- After a file is written, do not rewrite it for a few minutes.

- Delete and recreate files instead of overwriting the same file repeatedly.
- Use Memory Channel cluster interconnect.

To avoid overwhelming the duplicate request cache:

- Do not run hundreds of simultaneous processes that write files
- Do not operate the system under so heavy a load that NFS operations frequently take several seconds to complete.

Use the `netstat` command to determine if your network is saturated. For Ethernet networks, a high number of collisions indicates that the network may be saturated. The following example shows the output from the `netstat -I tu0` command:

Name	Mtu	Network	Address	Ipkts	Ierrs	Opkts	Oerrs	*Coll*
tu0	1500	<Link>	xx:xx:xx:xx:xx	840386045	0	254319298	5121	5014223
tu0	1500	network	client	840386045	0	254319298	5121	5014223
tu0	1500	DLI	none	840386045	0	254319298	5121	5014223

1.1.4 Tuning the NFS Server Duplicate Request Cache (new)

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

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

Patch 1062.00 provides a tuning variable to control the size of the NFS server's duplicate request cache:

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

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

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

1.1.5 Moving Files in DMAPI-Enabled Filesets Across Partitions — Patch 2084.00 (new)

When moving files in DMAPI-enabled filesets across partitions, the `mv` command will preserve only the default and access control list (ACL) attributes.

1.1.6 Performance of hwmgr Commands on Large System Configurations (new)

On large system configurations, certain `hwmgr` commands may take a long time to run and can produce voluminous output. For example, on a system connected to a large storage area network, the `hwmgr -view devices` command can take

a long time to begin displaying output, because it must first select devices from all of the hardware components in the system and then retrieve, format, and sort the output report.

The output from these commands is gathered and sorted in memory before the report begins to be displayed. In a system with hundreds or thousands of attached storage units, the processing stage can take several minutes and you will not see any output during that time.

When using the command `hwmgr -view devices -cluster`, the time can be even longer and the size of the report can be larger because in most clustered configurations, mass storage devices are reported by every member and thus appear multiple times in the generated report. Therefore, you may need to relax the memory limits for the process running the command, because with such a large number of devices in the configuration, the system may be unable to gather all of the data with the default memory limit.

We recommend that you run commands that generate large reports in the background (for example, in a batch job) and save their output into a file or set of files for subsequent examination or historical comparison.

1.1.7 Packetfilter Enhancements — Patch 2084.00 (new)

Patch 2084.00 provides enhancements with the following dbx kernel flags to control packetfilter-written packets:

- `pfilt_loopback=[0|1]`

Setting this flag to 0 prevents a loop back of any packetfilter-written multicast or broadcast packet.

When set to 1, the packetfilter loops back both broadcast and multicast packets. This is the default

- `pfilt_physaddr=[0|1]`

Setting this flag to 0 allows the application to fill in the source Ethernet address for packetfilter-written packets. If the source address is all 0's, the address is set to the proper hardware address by the packetfilter code as a safety precaution.

When set to 1, the packetfilter sets the Ethernet source address in the outgoing packet. This is the default.

Restriction

The `pfilt_loopback` and `pfilt_physaddr` tunable variables are only accessible in the kernel by using dbx.

In a future patch kit, these tunable variables will be implemented as kernel subsystem configurable attributes, accessible through the `sysconfig` command.

1.1.8 Changes to `vdump` and `vrestore` Allow Larger Record Sizes — Patch 2064.00

The `vdump` and `vrestore` programs have been tuned to disallow block sizes greater than 64 KB blocks. This is to avoid forward compatibility problems. With this patch, the valid range is 2 through 64 KB blocks.

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

The vrestore program still has the capability to autosize vdump archives. Also, it is backward compatible, which means it can restore archives created by earlier versions of vdump.

Note

If you create an archive using this version of vdump, uninstalling Patch 2064.00 will prevent you from restoring that archive. This occurs because removing Patch 2064.00 would revert vdump and vrestore to earlier versions, which do not recognize archives created by the later versions. If this occurs, you will have to reinstall Patch 2064.00 to restore the archive.

1.1.9 Potential Problem Patching a System with WLS Subsets Installed

If you will be installing or removing patches on a system with Worldwide Language Subsets installed, you may encounter a shell limitation on the size of the shell environment. This is due to the large number of installed subsets. If you encounter this situation, your system may hang or you may see an error message similar to the following:

```
/usr/sbin/setld: arg list too long.
```

To prevent this from occurring, enter the following command to relax the command-line limits prior to running dupatch:

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

Note

Do not use this kernel setting as a default; enable it only when you encounter a problem where the `exec()` argument size limit has been approached.

This vulnerability will be fixed in a future patch kit.

1.1.10 Security Feature — Patch 2084.00

Patch 2084.00 provides a new security feature to prevent the execution of instructions that reside in heap or other data areas of process memory. The result is additional protection against buffer overflow exploits. This feature is similar in concept to Tru64 UNIX executable stack protection.

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

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

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

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

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

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

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

Note

The Java language interprets bytecode at run time. Unless marked as exempt, privileged applications written in Java will receive an error when they attempt to execute instructions residing in the unexecutable memory. The manner in which these errors are handled is application specific and thus unpredictable. This is why you must run the `/usr/sbin/javaexecutedata` before you enable `executable_data`.

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

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

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

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

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

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

For more information see `chatr(1)`.

1.1.11 Changes to tar, pax, and cpio Behavior — Patch 2088.00

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

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

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

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

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

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

1.1.12 Ignore Special Instruction for Patch 1924.00

When installing this kit, the `dupatch` utility incorrectly indicates a Special Instruction for Patch 1924.00. You can safely disregard this notice.

1.1.13 Russian Keyboard — Patch 1197.00

The 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 ]
};
```

1.1.14 Problem Seen on Systems with Smart Array Controller

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

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

If this occurs, take the following steps:

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

```
ciiss:
ciiss_throttle_threshold=5
```

2. Execute the following command:

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

3. Reboot your system:

```
# shutdown -r now
```

1.1.15 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.

1.1.16 Updates to sys_check

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

1.1.16.1 TMPDIR Variable

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

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

1.1.16.2 sys_check Version 125 Web Kit

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

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

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

1.1.17 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/dds_config (see dds_config(8) for more information).

1.1.18 New libots3 Library — 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
```

1.1.19 New esmd Daemon — 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.

1.1.20 Command Updates and Other Changes — Patch 2084.00

The following sections describe updates to commands delivered in Patch 2084.00.

1.1.20.1 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

1.1.20.2 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.

1.1.20.3 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.

1.1.20.4 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.

1.1.20.5 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.

1.1.20.6 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]

1.1.20.7 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.

1.1.20.8 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. Patch 1717.00 updates the `sys_attrs_ee(5)` reference page.

1.1.20.9 Support for Network Link Aggregation

This patch enables support for network link aggregation, which can 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. Patch 1651.00 updates the `lag(7)` reference page and Patch 1755.00 updates the `lagconfig(8)` reference page.

1.2 Summary of Base Operating System Patches

This section provides brief descriptions of the patches in Patch Kit 6 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 5)

Indicates a patch that was new in the indicated Version 5.1A patch kit.

- Existing

Indicates a patch that was introduced in Patch Kit 1 or Patch Kit 2.

- 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

- 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 86.00

Abstract: Fix for Korn shell hang

State: Existing

- 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 108.00

Abstract: Fixes a potential race deadlock

State: Existing

- 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

- 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

- Corrects a memory leak in the XTI socket code.

Number: Patch 143.00

Abstract: Allows `dxsetacl` utility to delete access ACLs

State: Existing

- Allows the `dxsetacl` utility to delete access ACLs.
-

Number: Patch 145.00

Abstract: Security (SSRT0638U)

State: Existing

- Corrects 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

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

Number: Patch 156.00

Abstract: Fixes problems which prevented envmond from starting

State: Existing

- Fixes problems which prevented envmond from starting.

Number: Patch 169.00

Abstract: Fixes a problem in latsetup

State: Existing

- 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

- 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

- 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

- Corrects a problem in the rdist utility which was causing segmentation faults on files with more than one link.

Number: Patch 189.00

Abstract: Fix for no rerouting problem on a CFS server

State: Existing

- 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 210.00

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

State: Existing

- 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

- Fixes a problem of the ATM setup script failing when configuring an elan if the lane subsystem is not loaded.

Number: Patch 224.00

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

State: Existing

- Fixes a problem where join may fail to clean up its lock files in /var/join.

Number: Patch 238.00

Abstract: Fix for dxsetacl utility

State: Existing

- Allows the dxsetacl utility to delete access ACLs.

Number: 245.00

Abstract: Fixes a problem in the strtod routine

State: Existing

- 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)

- Provides enablers for the Compaq Database Utility.

Number: Patch 259.00

Abstract: Removes extraneous header comments

State: Existing

- Removes extraneous history edit comments from exported DECthreads header files.

Number: Patch 281.00

Abstract: Fix for NHD kit installations

State: Existing

- 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

- 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. Subsequent attempts result in error messages that report on the existence of duplicate host names.

Number: Patch 327.00

Abstract: Fixes C++ incompatibility

State: Existing

- 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

- 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

- 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

- 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 446.00

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

State: Existing (Kit 4)

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

Number: Patch 453.00

Abstract: Fix for dtgreet application

State: Existing

- 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

- 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

- 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

- 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

- 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

- Fixes a problem in which a beep does not occur when requested when the toggle keys option is enabled via accessx.

Number: Patch 481.00

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

State: Existing

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

Number: Patch 485.00

Abstract: Fix for C++ compile problem

State: Existing

- Fixes a C++ compile problem in /usr/include/X11/Xlib.h.

Number: Patch 490.00

Abstract: Fix for class scheduler failure

State: Supersedes Patch 183.00

- Fixes a class scheduler semaphore race condition.
- Causes the automatic detection of 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

- Fixes an libXm.so incompatibility.

Number: Patch 519.00

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

State: Existing

- Fixes the C++ incompatibility of /usr/include/dec/pwrmgr/pwrmgr.h.

Number: Patch 525.00

Abstract: Security (SSRT0779U)

State: Existing

- Corrects a potential security vulnerability where, under certain circumstances, SNMP services can stop functioning.

Number: Patch 527.00

Abstract: Security (SSRT0779U)

State: Existing

- Corrects 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)

- 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

- 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

- Fixes problem in which the balance utility terminated before balancing the whole domain when the domain was very large (greater than 4 GB).

Number: Patch 537.00

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

State: Existing (Kit 3)

- 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

- 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

- Fixes an ATM signaling problem.

Number: Patch 553.00

Abstract: EVM daemon fails to find user-defined templates

State: Existing

- 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

- Provides enablers for the Compaq Database Utility.

Number: Patch 569.00

Abstract: CD Mastering Software

State: Existing

- 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

- 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

- 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

- 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

- 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)

- 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.

Number: Patch 693.00

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

State: Existing (Kit 3)

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

Number: Patch 701.00

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

State: Existing (Kit 3)

- 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)

- 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

- Clarifies the output of shfragbf, an AdvFS utility.
-

Number: Patch 713.00

Abstract: Fix for rcinet script

State: Existing

- 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 731.00

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

State: Existing (Kit 3)

- 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)

- Addresses a situation in which the Memory Channel device shuts down if too many state change interrupts are received.

Number: Patch 741.00

Abstract: Enhancement to savemeta script

State: Existing (Kit 3)

- 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)

- 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)

- 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 751.00

Abstract: Security (SSRT0818U)

State: Existing (Kit 3)

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

Number: Patch 753.00

Abstract: Security (SSRT0818U)

State: Existing (Kit 3)

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

Number: Patch 755.00

Abstract: Security (SSRT0818U)

State: Existing (Kit 3)

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

Number: Patch 757.00

Abstract: Fix for script command

State: Existing (Kit 3)

- 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)

- 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)

- 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)

- 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)

- Adds the poll(7) reference page for the /dev/poll driver.

Number: Patch 767.00

Abstract: Enhancement to fuser utility

State: Existing

- 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)

- 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)

- 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) • Corrects 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) • Fixes an xfs problem which causes a "QueryGlyphs failed" error in showfont.
Number: Patch 789.00 Abstract: Fix for ppdfof print filter core dump problem State: Existing (Kit 3) • 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) • 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) • 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) • 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) • 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 809.00 Abstract: Enables correctable error reporting from DTAGII State: Existing (Kit 3) • 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: Supersedes Patches 165.00, 167.00, 557.00 • 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)

- Delivers the ckfsec(1) reference page.

Number: Patch 826.00

Abstract: Security (SSRT0794U)

State: Existing (Kit 3)

- Corrects 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)

- Corrects 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.
- 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 of SysMan tools on the iPAQ and those 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)

- 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)

- 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: Existing (Kit 4)

- 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: Existing (Kit 4)

- 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

- Corrects 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

- 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: Existing (Kit 4)

- Fixes many small problems in dsfmgr.

Number: Patch 999.00

Abstract: Security (SSRT1-80U)

State: Supersedes Patches 71.00 and 997.00

- 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

- Corrects 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 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

- 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 1027.00

Abstract: Security

State: Existing (Kit 4)

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

Number: Patch 1044.00

Abstract: Fixes failures under high DMA resource utilization

State: Existing (Kit 4)

- 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: Existing (Kit 4)

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

Number: Patch 1061.00

Abstract: Security (SSRT2208)

State: Existing (Kit 4)

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

Number: Patch 1063.00

Abstract: Fixes hwmgr command to show path state correctly

State: Existing (Kit 4)

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

Number: Patch 1073.00

Abstract: Security (SSRT2229)

State: Existing (Kit 4)

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

Number: Patch 1075.00

Abstract: Security

State: Existing (Kit 4)

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

Number: Patch 1087.00

Abstract: Security

State: Existing (Kit 4)

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

Number: Patch 1089.00

Abstract: Security

State: Existing (Kit 4)

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

Number: Patch 1091.00

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

State: Existing (Kit 4)

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

Number: Patch 1093.00

Abstract: Security (SSRT2339, SSRT2339)

State: Existing (Kit 4)

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

Number: Patch 1095.00

Abstract: Security

State: Existing (Kit 4)

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

Number: Patch 1097.00

Abstract: Fix prevents "simple lock owned" panics

State: Existing (Kit 4)

- Prevents "simple lock owned" panics.

Number: Patch 1099.00

Abstract: Security

State: Existing (Kit 4)

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

Number: Patch 1101.00

Abstract: Security

State: Existing (Kit 4)

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

Number: Patch 1103.00

Abstract: Fix for verify command

State: Existing (Kit 4)

- 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

- 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 1109.00

Abstract: Security

State: Existing (Kit 4)

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

Number: Patch 1115.00

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

State: Existing (Kit 4)

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

Number: Patch 1117.00

Abstract: Correct improper file or privilege management

State: Existing

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

Number: Patch 1121.00

State: Existing (Kit 4)

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

Number: Patch 1123.00

Abstract: Security

State: Existing (Kit 4)

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

Number: Patch 1125.00

Abstract: Fix for hwmgr -show name command

State: Existing (Kit 4)

- Fixes the display for the hwmgr -show name command to align properly for the name field.
-

Number: Patch 1127.00

Abstract: Security

State: Existing (Kit 4)

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

Number: Patch 1129.00

Abstract: Update to hwmgr utility

State: Existing (Kit 4)

- 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 1140.00

Abstract: Modifies enablers for Enterprise Volume Manager

State: Supersedes Patches 199.00, 267.00, 315.00

- Modifies enablers for the Enterprise Volume Manager.

Number: Patch 1142.00

Abstract: Security (SSRT0792U)

State: Existing (Kit 4)

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

Number: Patch 1144.00

Abstract: Security

State: Existing (Kit 4)

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

Number: Patch 1146.00

Abstract: Installs version V2.1-120 of libots3

State: Supersedes Patch 226.00

- 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

- 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

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

Number: Patch 1156.00

Abstract: Security

State: Existing (Kit 4)

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

Number: Patch 1166.00

Abstract: Fixes consvar -s bootdef_dev failure with KZPCC

State: Existing (Kit 4)

- Fixes a problem where I/O greater than 4 MB 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

- 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 1172.00

Abstract: Fix for sysconfig utility

State: Existing (Kit 4)

- Fixes a problem in which the lines in the output stream from sysconfig -Q can be truncated.

Number: Patch 1175.00

Abstract: Security

State: Existing (Kit 4)

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

Number: Patch 1183.00

Abstract: Security

State: Existing (Kit 4)

- Corrects 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: Existing (Kit 4)

- Corrects 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: Existing (Kit 4)

- Corrects 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 1192.00

Abstract: Security

State: Existing (Kit 4)

- Corrects 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: Existing (Kit 4)

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

Number: Patch 1202.00

Abstract: Security

State: Existing (Kit 4)

- Corrects 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: Existing (Kit 4)

- 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

- 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: Existing (Kit 4)

- Fixes a simple lock panic in the floppy driver.

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

- 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 1224.00

Abstract: Security

State: Existing (Kit 4)

- Corrects 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 1228.00

Abstract: Correct improper file or privilege management

State: Existing

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

Number: Patch 1232.00

Abstract: Security (SSRT2368, SSRT2368)

State: Existing (Kit 4)

- Corrects 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 1242.00

Abstract: libpset shared library fix

State: Existing (Kit 4)

- 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: Existing (Kit 4)

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

Number: Patch 1246.00

Abstract: Security

State: Existing (Kit 4)

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

Number: Patch 1253.00

Abstract: Security

State: Existing (Kit 4)

- Corrects 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: Existing (Kit 4)

- Corrects 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: Existing (Kit 4)

- Corrects 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

- 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: Existing (Kit 5)

- Corrects 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: Existing (Kit 4)

- 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 helps distinguish between CPU bindings and RAD bindings, and also keeps 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 1274.00

Abstract: Security

State: Existing (Kit 4)

- Corrects 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

- 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 1281.00

Abstract: Corrects a problem with SNMP

State: Existing (Kit 4)

- 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 1293.00

Abstract: Security

State: Existing (Kit 4)

- Corrects 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: Existing (Kit 4)

- Corrects 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: Existing (Kit 4)

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

Number: Patch 1307.00

Abstract: Fixes a problem with the operation of osf_boot

State: Existing (Kit 4)

- Fixes a problem with the operation of osf_boot with the interactive flag on an AlphaServer ES45 system.

Number: Patch 1309.00

Abstract: Security

State: Existing (Kit 4)

- Corrects 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: Existing (Kit 4)

- Corrects 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: Existing (Kit 4)

- Corrects 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 1319.00

Abstract: Security

State: Existing (Kit 4)

- Corrects 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 1332.00

Abstract: Security

State: Existing (Kit 4)

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

Number: Patch 1334.00

Abstract: Security

State: Supersedes Patches 193.00 and 1119.00

- 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: Existing (Kit 4)

- 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 1350.00

Abstract: Security (SSRT2193)

State: Existing (Kit 4)

- 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 1359.00

Abstract: Corrects hang in the log command

State: Existing

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

Number: Patch 1370.00

Abstract: Revises the ifconfig.8.gz reference page

State: Existing (Kit 5)

- Revises the ifconfig(8) reference page.

Number: Patch 1372.00

Abstract: Updates for the zoneinfo data file

State: Existing (Kit 5)

- Updates the time zone data files in /etc/zoneinfo/ to include the most recent changes from the latest PD time zone source kit (tzdata2002d).

Number: Patch 1374.00

Abstract: Corrects improper file or privilege management

State: Existing (Kit 5)

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

Number: Patch 1376.00

Abstract: Blocked mutex lock causes XTI problem

State: Existing (Kit 5)

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

Number: Patch 1378.00

Abstract: Threaded apps using XTI/TLI may terminate or hang

State: Existing (Kit 5)

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

Number: Patch 1539.00

Abstract: Scripts in /sbin/init.d are now world-readable

State: Existing (Kit 5)

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

Number: Patch 1541.00

Abstract: Security (SSRT2275)

State: Supersedes Patches 1018.00, 1020.00, 469.00

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file or privilege management.
- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised when a buffer overflow occurs in the uucp utility. Buffer overflows are sometimes exploited in an attempt to subvert the function of a privileged program and possibly execute commands at the elevated privileges if the program file has the setuid privilege.
- Fixes a problem in uucp. uucp between two Tru64 UNIX boxes hangs when a uucp failure occurs.
- Provides protection against a class of potential security vulnerabilities called buffer overflows. Buffer overflows are sometimes exploited in an attempt to subvert the function of a privileged program and possibly execute commands at the elevated privileges if the program file has the setuid privilege. This patch allows a system administrator to enable memory management protections that limit potential buffer overflow vulnerabilities.

Number: Patch 1543.00

Abstract: Security (SSRT2275)

State: Supersedes Patch 1266.00

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

Number: Patch 1545.00

Abstract: Scripts in /sbin/init.d are now world-readable

State: Existing (Kit 5)

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

Number: Patch 1547.00

Abstract: Security (SSRT2275)

State: Existing (Kit 5)

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

Number: Patch 1551.00

Abstract: Scripts in /sbin/init.d are now world-readable

State: Existing (Kit 5)

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

Number: Patch 1553.00

Abstract: Scripts in /sbin/init.d are now world-readable

State: Existing (Kit 5)

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

Number: Patch 1555.00

Abstract: Fixes various problems in the libc functions

State: Existing (Kit 5)

- Fixes various problems in the libc functions getdate(), strptime(), callrpc(), strncasecmp(), fork()/popen(), and ncreate().

Number: Patch 1565.00

Abstract: Security (SSRT0753U, SSRT0752U, SSRT0788U)

State: Supersedes Patches 433.00, 434.00, 436.00, 697.00, 1009.00, 1011.00

- Fixes the dtprintinfo memory fault problem with long LANG value.
- Fixes several potential security vulnerabilities where, under certain circumstances, system integrity may be compromised. This may be in the form of large values of command line arguments.
- Corrects 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.
- Corrects a potential security vulnerability in CDE Subprocess Control Service (dtspcd), which has a potential buffer overflow condition that may lead to 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 access.
- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised when a buffer overflow occurs in the DtSvc utility. Buffer overflows are sometimes exploited in an attempt to subvert the function of a privileged program and possibly execute commands at the elevated privileges if the program file has the setuid privilege.

Number: Patch 1567.00

Abstract: tcpdump does not filter UDP traffic properly

State: Existing (Kit 5)

- Fixes a problem in tcpdump that caused it to not filter UDP traffic properly.
-

Number: Patch 1570.00

Abstract: Provides enhanced KDBX debugging features

State: Supersedes Patches 699.00, 1353.00, 1568.00

- Fixes a premature termination of the ofile kdbx extension, warning messages in various kdbx extensions, and token length warnings when kdbx is invoked.
- Fixes a problem with audit data not being displayed by audit tool, problems with file object selection/deselection and directories, and NUMA performance issues associated with auditing.
- Fixes problems in the kdbx u and vnode extensions.
- Enhances KDBX debugging features to include a -A flag for route and to keep inpcb from truncating port numbers.

Number: Patch 1572.00

Abstract: Fix for termcap script

State: Existing (Kit 5)

- Fixes a problem during a cluster rolling upgrade in which the merge for the termcap file failed.

Number: Patch 1574.00

Abstract: Updates to Sysman Account Management application

State: Supersedes Patches 1057.00, 1059.00

- Corrects a problem in which errors are generated when nonroot users who lack a valid home directory run the SysMan Account Management tool.
- Corrects a problem in which errors are generated when +users or -users are present in the /etc/passwd file and the SysMan Account Management tool is used to display or modify +user or -user. The symptom is: "several property fields, such as comments, LoginShell, HomeDirectory, for the +user or -user will have values of 1 or 0".
- Fixes a problem in dealing with white spaces in a "Filter by Comment" search of the SysMan Account Management application.

Number: Patch 1587.00

Abstract: Correct improper file or privilege management

State: Supersedes Patches 1113.00, 1584.00, 1585.00

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file or privilege management.
- Allows accounting files to be referenced using CDSLs.
- Corrects the display of the header from acctcom when accounting is first started.
- Fixes an error in the lastlogin.sh script.
- Resolves differences in CPU and connect times found in the conversion from ASCII format to binary and back to ASCII in accounting reports.
- Resolves the differences in CPU time found in the output of the acctcom and acctmerg commands for the same input file.

Number: Patch 1589.00

Abstract: fwtmp will not display invalid (negative) pids

State: Existing (Kit 5)

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

Number: Patch 1591.00

Abstract: Removes the 250 variable limit for env command

State: Existing (Kit 5)

- Removes the 250 variable limit for /usr/bin/env.
-

Number: Patch 1593.00

Abstract: Security (SSRT2408, SSRT2411, SSRT2410)

State: Existing (Kit 5)

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

SSRT2408 BIND - (Severity - High)

SSRT2410 BIND - (Severity - High)

SSRT2411 BIND - (Severity - High)

Number: Patch 1596.00

Abstract: Corrects improper file or privilege management

State: Existing (Kit 5)

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

Number: Patch 1604.00

Abstract: mcopy/mwrite can now overwrite existing files

State: Existing (Kit 5)

- Eliminates several security vulnerabilities.
 - Changes mcopy and mwrite to let them overwrite existing files.
 - Fixes mtools to let it return appropriate error messages for nonprivileged users.
 - Standardizes the mformat prompt.
-

Number: Patch 1609.00

Abstract: Correct improper file access

State: Supersedes Patch 1111.00

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

Number: Patch 1611.00

Abstract: Revises the kdbx(8) reference page

State: Existing (Kit 5)

- Revises the kdbx(8) reference page to document the addition of the -A flag to route.
-

Number: Patch 1613.00

Abstract: Fix for cdvd command

State: Existing (Kit 5)

- Corrects the output of the CDVD command, which had one extraneous line and a line that was incorrectly labeled.
-

Number: Patch 1618.00

Abstract: Fixes message catalog problem in rrestore

State: Supersedes Patches 506.00, 1616.00

- Corrects the rdump command to dump data properly onto remote tape devices without receiving the SIGSEGV and dumping core.
 - Fixes dump to recognize LSM volumes correctly and not report random information when an error has occurred.
 - Introduces dumpprmt.msg for remote dump and restore messages. This new message catalog file is used in both rdump and rrestore programs.
-

Number: Patch 1620.00

Abstract: Eliminates compiler warnings in mkdir

State: Existing (Kit 5)

- Eliminates compiler warnings in mkdir.

Number: Patch 1622.00

Abstract: Update to adduser command

State: Existing (Kit 5)

- Corrects the default UID displayed by adduser when UID_MAX exists in the /etc/passwd file.
- Causes adduser to avoid duplicating UIDs.

Number: Patch 1624.00

Abstract: Adds the chatr(1) reference page

State: Existing (Kit 5)

- Adds the chatr(1) reference page.

Number: Patch 1626.00

Abstract: Enhancement to chfile command

State: Existing (Kit 5)

- Causes the chfile command to display a more informative error message if chfile fails while trying to enable data logging.

Number: Patch 1628.00

Abstract: Enables tip to log into member specific log file

State: Existing (Kit 5)

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

Number: Patch 1630.00

Abstract: Path for the aculog file has been corrected

State: Existing (Kit 5)

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

Number: Patch 1632.00

Abstract: Provides correct labels for audit_tool and auditmask

State: Existing (Kit 5)

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

Number: Patch 1634.00

Abstract: Audit subsystem utilities printing out wrong labels

State: Existing (Kit 5)

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

Number: Patch 1636.00

Abstract: Correct improper file or privilege management

State: Existing (Kit 5)

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

Number: Patch 1638.00

Abstract: Corrects improper file access

State: Existing (Kit 5)

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

Number: Patch 1641.00

Abstract: Fixes race condition in rm command

State: Existing (Kit 5)

- Addresses the performance issue of rm -r with large directories.
- Fixes the problem of a race condition in rm command, wherein two threads can successfully delete a file simultaneously.

Number: Patch 1647.00

Abstract: quotacheck may incorrectly report disk quote excess

State: Existing (Kit 5)

- Corrects a problem where the quotacheck utility may incorrectly report that a disk quota has been exceeded.

Number: Patch 1649.00

Abstract: Problem with execution of bttape TCL scripts

State: Existing (Kit 5)

- Fixes a problem encountered when bttape TCL scripts are executed by nonprivileged users.
- Adds \$quote directive to the message catalog.

Number: Patch 1651.00

Abstract: Revises the lag(7) reference page

State: Existing (Kit 5)

- Revises the lag(7) reference page.
- Updates information about ee drivers for DE60x Ethernet cards.

Number: Patch 1653.00

Abstract: Fixes a problem in the KZPCA itpsa driver

State: Supersedes Patches 261.00,483.00, 797.00, 1233.00, 1235.00

- 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 a SDLT media error that causes bus resets with KZPCA adapters.
- Fixes a problem in the KZPCA itpsa driver that can be seen when a SCSI target presents multiple LUNs.

Number: Patch 1665.00

Abstract: Corrects the behavior of ln -sf

State: Existing (Kit 5)

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

Number: Patch 1667.00

Abstract: Kernel memory corruption in subscription routine

State: Existing (Kit 5)

- Corrects a possible kernel memory corruption problem in the kernel event subscription routine.
-

Number: Patch 1669.00

Abstract: Revises the sys_attrs_dli(5) reference page

State: Existing (Kit 5)

- Revises the sys_attrs_dli(5) reference page to document the two new globals dli_sendspace and dli_recvspace.

Number: Patch 1675.00

Abstract: Correct improper file or privilege management

State: Supersedes Patch 1272.00

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file or privilege management.
- Corrects the behavior of more, when given both a non-existing file and a non-empty file with long a file name or pathname.

Number: Patch 1677.00

Abstract: advfsstat prints negative values for statistics

State: Existing (Kit 5)

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

Number: Patch 1679.00

Abstract: cut command now handles incomplete lines correctly

State: Existing (Kit 5)

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

Number: Patch 1682.00

Abstract: Security (SSRT0743U, SSRT2256)

State: Supersedes Patches 69.00, 1177.00, 1680.00

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file or privilege management.
- Corrects a 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.
- Fixes a situation, in which the header fields displayed by the output of ps command are not aligned properly.
- Allows white space in header field with ps -o. Multiple headers with white space can be given with ps -o.

Number: Patch 1684.00

Abstract: Corrects label strings in sysman LSM application

State: Existing (Kit 5)

- Corrects some label strings in the SysMan LSM application.

Number: Patch 1686.00

Abstract: S19security script causes change in security config

State: Existing (Kit 5)

- Corrects a problem in which running the /sbin/rc2.d/S19security script during system start up could cause an unintended change in the system security configuration. This would happen only when /usr/bin/perl had been removed.

Number: Patch 1688.00

Abstract: Buffer overflow problem in the write command

State: Existing (Kit 5)

- Fixes a buffer overflow problem in /usr/bin/write.
-

Number: Patch 1693.00

Abstract: Corrects improper file or privilege management

State: Existing (Kit 5)

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

Number: Patch 1695.00

Abstract: grep now allow blank lines in the pattern file

State: Existing (Kit 5)

- Changes the grep command to now allow blank lines in the pattern file and does not hang when executed with -w and -f options.

Number: Patch 1697.00

Abstract: defragment fails when unable to obtain enough memory

State: Existing (Kit 5)

- Fixes a problem where defragment can fail if it is unable to obtain enough memory.

Number: Patch 1699.00

Abstract: Fix for volmake utility

State: Existing (Kit 5)

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

Number: Patch 1701.00

Abstract: Security (SSRT2275)

State: Existing (Kit 5)

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

Number: Patch 1703.00

Abstract: X server may hang every 49 days with Powerstorm card

State: Supersedes Patch 1042.00

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

Number: Patch 1707.00

Abstract: Modifies enablers for Enterprise Volume Manager

State: Supersedes Patches 197.00, 263.00, 313.00, 513.00, 1136.00

- Modifies enablers for the Enterprise Volume Manager.

Number: Patch 1709.00

Abstract: Modifies enablers for Enterprise Volume Manager

State: Supersedes Patches 201.00, 265.00, 317.00, 515.00, 1138.00

- Modifies enablers for the Enterprise Volume Manager.

Number: Patch 1711.00

Abstract: Removes compiler warnings addressing outside of array

State: Existing (Kit 5)

- Removes compiler warnings addressing outside of array bounds.
-

Number: Patch 1713.00

Abstract: Fixes a problem with os_mibs

State: Supersedes Patch 739.00

- Fixes a problem with os_mibs that could cause the application to consume an excessive amount of CPU time.
- Corrects a problem in os_mibs which resulted in the swap size and swap used values for the host mib being reported as negative values on some systems.

Number: Patch 1715.00

Abstract: Corrects exit status of sed when disk is full

State: Existing (Kit 5)

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

Number: Patch 1717.00

Abstract: Revises the sys_attrs_ee(5) reference page

State: Existing (Kit 5)

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

Number: Patch 1719.00

Abstract: Corrects improper file access

State: Supersedes Patches 529.00, 1158.00

- Fixes a problem where no shell message is displayed when trying to su to a user other than root.
- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file access.
- Addresses problems with the mksas utility where false warning messages are generated and where the user-specified temporary directory could be erroneously removed.

Number: Patch 1721.00

Abstract: Corrects improper file or privilege management

State: Existing (Kit 5)

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

Number: Patch 1729.00

Abstract: Corrects which to take path info from environment

State: Existing (Kit 5)

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

Number: Patch 1731.00

Abstract: Corrects improper file or privilege management

State: Existing (Kit 5)

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

Number: Patch 1735.00

Abstract: Corrects problems with the NIFF daemon

State: Supersedes Patch 220.00

- Corrects a problem where the NIFF daemon (nifd) would exit if its connection to the EVM daemon (evmd) failed, as in the case of an EVM daemon restart.
 - Corrects a problem in nifd that results in its memory usage growing over time.
-

Number: Patch 1737.00

Abstract: Fix for bextract command

State: Existing (Kit 5)

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

Number: Patch 1739.00

Abstract: make command not checking for time stamps

State: Existing (Kit 5)

- Changes /usr/opt/ultrix/usr/bin/make so that it properly checks dependencies on archive libraries.

Number: Patch 1741.00

Abstract: Fix for login script

State: Existing (Kit 5)

- Corrects a problem during a rolling upgrade in which the merge of the .login file would fail, with inadequate informational messages.

Number: Patch 1743.00

Abstract: Order of records in CDF file not preserved

State: Existing (Kit 5)

- Corrects a problem in which restoring cloned information from a generated configuration description file (CDF) using sysman -clone -apply <file> fails to preserve the order of records in the CDF file. CDF files are used to clone systems at the configuration and installation level.

Number: Patch 1745.00

Abstract: Fixes binlogd core dump problem

State: Supersedes Patches 586.00, 777.00, 1214.00

- 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.
- Corrects a potential binlogd core dump problem when parsing its remote host authorization file (/etc/binlog.auth) with greater than 513 characters.

Number: Patch 1747.00

Abstract: Corrects improper file or privilege management

State: Existing (Kit 5)

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

Number: Patch 1749.00

Abstract: SysMan Station may fail to generate hardware view

State: Existing (Kit 5)

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

Number: Patch 1751.00

Abstract: Corrects improper file or privilege management

State: Existing (Kit 5)

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

Number: Patch 1753.00

Abstract: Enhancement for siacfg utility

State: Supersedes Patch 787.00

- Allows non-local SIA mechanisms, LDAP for example, to place their mechanism last in the list of mechanisms.
- Eliminates the "Using an array as a reference is deprecated" warning when running /usr/sbin/siacfg and during system boot on systems using Perl 5.8.0 and higher.

Number: Patch 1755.00

Abstract: Revises the lagconfig(8) reference page

State: Existing (Kit 5)

- Revises the lagconfig(8) reference page to document fixes for Link Aggregation.

Number: Patch 1757.00

Abstract: Corrects an error return code for volinfo utility

State: Existing (Kit 5)

- Corrects an error return code for volinfo. If a specified volume is not in the configuration database, the error code returned from volinfo will reflect this error.

Number: Patch 1759.00

Abstract: mkpasswd command dumps core

State: Existing (Kit 5)

- Fixes a situation in which the mkpasswd command dumps core when executed by a nonprivileged user.

Number: Patch 1761.00

Abstract: Corrects improper file or privilege management

State: Existing (Kit 5)

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

Number: Patch 1763.00

Abstract: Fix for bcheckrc script

State: Supersedes Patch 1278.00

- Clarifies a vague message that dn_setup would print if a system was brought to single-user mode because dsfingr detected database errors during boot.
- Fixes a problem with bcheckrc that occurs when it is run multiple times.

Number: Patch 1765.00

Abstract: Corrects improper file or privilege management

State: Existing (Kit 5)

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

Number: Patch 1767.00

Abstract: Fixes IPv6 neighbor discovery daemon problem

State: Existing (Kit 5)

- Fixes a problem where, under cerntain circumstances, the IPv6 neighbor discovery daemon can cause bad information to be written to a DNS database causing failures on subsequent database reloads.

Number: Patch 1769.00

Abstract: Fixes re_ioctl() cases DIODCMD and DIODCDB

State: Existing (Kit 5)

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

Number: Patch 1771.00

Abstract: Security (SSRT0664U, SSRT0762U, SSRT0801)

State: Supersedes Patches 204.00, 206.00, 567.00, 1303.00, 1305.00

- Corrects a potential security vulnerability 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 that 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 potential security vulnerability which, under certain circumstances, could compromise system integrity.
- Corrects a problem in the ftp open command to allow the optional port argument to accept port numbers between 32768 and 65535.
- Fixes an ftp client accounting error that occurs when transferring more than 4 GB files.

Number: Patch 1773.00

Abstract: scu utility displays misleading data expected pattern

State: Supersedes Patch 234.00

- Updates /sbin/scu, the SCSI CAM utility program, to add support for Persistent Reserve for HSV110 and the display of 128-bit WWIDS.
- Fixes a problem with scu where a mismatch between expected and found data displays incorrect data expected.

Number: Patch 1775.00

Abstract: Eliminates use of /tmp file in SysMan CLI example

State: New

- Eliminates the use of a /tmp file in a SysMan CLI example.

Number: Patch 1783.00

Abstract: Fix for rpc.lockd

State: Existing (Kit 5)

- Fixes issues with rpc.lockd dealing with replies to message passing RPCs, requests from hosts with multiple IP addresses, and grant messages issued to down clients.

Number: Patch 1786.00

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

State: Existing (Kit 5)

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

Number: Patch 1804.00

Abstract: Fix for screend daemon

State: Existing (Kit 5)

- Fixes a potential problem in screend.
-

Number: Patch 1813.00

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, 1348.00, 1605.00, 1607.00

- Corrects lpd parent daemon problems when EVM is stopped and started.
- Slows down event storm from remote host sending bad protocol information.
- Fixes the following problems with the C shell command interpreter, csh:
 - Corrects the error message displayed when a nonroot user performs issues the ls command with wildcard characters on a directory having permission 700.
 - Corrects the error message displayed when nonomatch is set and a user 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 a problem with csh redirection while redirecting standard input and standard output of a command to a file exist in a home directory using tilde (~) operation.
- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised when a buffer overflow occurs in the lpq, lpr and lprm commands. Buffer overflows are sometimes exploited in an attempt to subvert the function of a privileged program and possibly execute commands and the elevated privileges if the program file has the setuid privilege.
- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised 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 and the elevated privileges if the program file has the setuid privilege.
- Corrects a problem where telnetd leaves an extra udp port open.
- Corrects an lpc regression in the lpc buffer overflow fix.
- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised when a buffer overflow occurs in the csh utility. Buffer overflows are sometimes exploited in an attempt to subvert the function of a privileged program and possibly execute commands and the elevated privileges if the program file has the setuid privilege.
- Corrects several potential security vulnerabilities where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file access.
- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised when a buffer overflow occurs in the telnetd daemon. Buffer overflows are sometimes exploited in an attempt to subvert the function of a privileged program and possibly execute commands and the elevated privileges if the program file has the setuid privilege.
- Fixes a number of problems in the lpd (line printer daemon) subsystem.
- Fixes lpd case sensitivity, for example, "node.domain" being treated the same as "Node.Domain."
- Fixes a problem that causes a segmentation fault. when dbx is analyzing a Fortran program.

Number: Patch 1817.00

Abstract: mkcdsl now properly deals with sticky bits on files

State: Existing (Kit 5)

- Fixes a problem with mkcdsl not carrying the sticky bit through.
-

Number: Patch 1821.00

Abstract: EVM returns DECEvent startup error in email msg

State: Existing (Kit 5)

- Corrects a problem in which protocol handshake error occurs when a high priority binlog event is being reported via a mail message. This occurs because EVM uses DECEvent to translate the binlog event, which requires the HOME environment variable to be set in order to startup.

Number: Patch 1825.00

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

State: Existing (Kit 5))

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

Number: Patch 1831.00

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, 1181.00

- 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 1918.00

Abstract: Corrects Enhanced Security C1crypt password change

State: Supersedes Patches 1559.00, 1800.00, 1916.00

- Fixes client (login, su, rshd, edauth, and sshd2) hangs and long delays under Enhanced Security, as well as some intermittent errors or failures seen with prpasswd or rpc.yppasswdd.
 - Corrects a problem in which logins in TruCluster environments using Enhanced Security could hang on any member other than the one serving /var to CFS.
 - Corrects a problem on systems running Enhanced Security in which the edauth -R command will not write user-profile entries to the root partition.
 - Corrects a problem in which using C1crypt for password encryption on Enhanced Security systems prevents users from changing their passwords and causes the passwd command state to display the message "Password not changed: failed to write protected password entry."
-

Number: Patch 1924.00

Abstract:Enabler for Enterprise Volume Manager product

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, 1471.00, 1451.00, 1516.00, 1498.00, 187.00, 1346.00, 1823.00, 517.00, 1673.00, 1919.00, 1920.00, 1921.00, 1922.00

- Proves enablers for the Enterprise Volume Manager product.
- Prevents a vold from core dumping when removing a disk from rootdg using voldiskadm or voldg.
- Prevents a KMF (kernel memory fault) panic, in voldiskstart(), when an I/O is attempted on an LSM device that is not accessible.
- Fixes a situation in which when a cluster member fails, mirrored volumes are left in a state such that recovery is always necessary when members boot, even if no additional recovery should be necessary.
- Prevents vold from core dumping on booting with the message:

Commit: not holding dg lock <diskgroup>

- 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 problems with the LSM commands volsave, volrestore and volclonedg.
 - Corrects performance issues on starting of Cluster Logical Storage Manager (CLSM) with large configurations.
 - Allows the proper synchronization of disk errors and failures in a cluster under CLSM control.
 - Fixes several potential security vulnerabilities where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file access.
 - Prevents inconsistent LSM volumes when the name of a partition that is being encapsulated matches the name of a current LSM volume.
 - 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. This patch prevents the smsd from triggering LSM configuration errors when querying LSM in a cluster.
 - Provides protection against a class of potential security vulnerabilities called buffer overflows. Buffer overflows are sometimes exploited in an attempt to subvert the function of a privileged program and possibly execute commands at the elevated privileges if the program file has the setuid privilege. This patch allows a system administrator to enable memory management protections that limit potential overflow vulnerabilities.
-

Patch 1924.00 Continued

- Corrects an I/O performance issue seen when CLSM is configured and one member of a cluster goes down unexpectedly.
- Corrects an issue of a cluster node hanging on boot while the other member recovers the cluster root file systems.
- Modifies volencap to prevent encapsulation of restricted partitions or placing swap into a non-rootdg diskgroup.
- Allow volsave and volrestore to save nconfig/nlog policies for diskgroups and to restore them appropriately.
- Corrects awk errors for invalid quit statements.
- 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.
- Fixes a problem in which volassist was unable to create a mirror of a striped volume with the mirror having a layout of concat.
- Improves null partition checking code.
- Fixes issues with I/O failures that are seen on non-rootdg diskgroups, including a conditions that causes the diskgroup to fail under certain conditions and to core dump when failing a non-rootdg disk.
- Fixes a performance issue seen during plex detach operations.
- Provides fixes that allow a full recovery in the event of an invalid magic number found in the recovery map.
- Modifies the volassist command to correctly output a warning message for its grow and shrink operations.
- Corrects an internationalization problem affecting the LSM voldiskadm command. An instructional message that is output when choosing the voldiskadm list option was incorrect because it had not been entered properly in the message catalog.

Number: Patch 1926.00**Abstract:** Security (SSRT2191)**State:**Supersedes 1317.00

- Corrects 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.
 - Improves null partition checking code.
-

Number: Patch 1928.00

Abstract: Security (SSRT2275)

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, 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, 1368.00, 1549.00, 1815.00

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file or privilege management.
- 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 it will not print "Succeeded" on a failure. For example:

showfdmn: unable to get info for domain 'domain_used' showfdmn: Succeeded

- Corrects a problem in which advscan command incorrectly processes concatenated options (such as -ar vs. -a -r). For example, if -ar is specified, the -r option is not processed.
-

Patch 1928.00 Continued

- Fixes a problem that prevents AdvFS from working correctly with LSM volumes between 1 TB and 2 TB.
- Adds more helpful informational messages, for example:
 - Advscan will now indicate 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 various minor problems in dsfmg.
- Fixes an rmvol E_PAGE_NOT_MAPPED error.
- Eliminate an ENO_MORE_BLK error seen when performing a copy-on-write procedure to a clone file while an rmvol operation is in progress.
- Provides protection against a class of potential security vulnerabilities called buffer overflows. Buffer overflows are sometimes exploited in an attempt to subvert the function of a privileged program and possibly execute commands at the elevated privileges if the program file has the setuid privilege. This patch allows a system administrator to enable memory management protections that limit potential buffer overflow vulnerabilities.
- Prevents erroneous DMAPI messages that were being printed while using showfile.
- Improves null partition checking code.

Number: Patch 1930.00**Abstract:** Correct edauth failure to write to root partition**State:** New

- Corrects a problem on systems running Enhanced Security in which the edauth -R command will not write user-profile entries to the root partition.

Number: Patch 1933.00**Abstract:** Corrects Enhanced Security C1crypt password change**State:** Supersedes 1557.00, 1798.00, 1931.00

- Fixes client (login, su, rshd, edauth, and sshd2) hangs and long delays under Enhanced Security, as well as some intermittent errors or failures seen with prpasswd or rpc.yppasswdd.
- Corrects a problem in which logins in TruCluster environments using Enhanced Security could hang on any member other than the one serving /var to CFS.
- Corrects a problem on systems running Enhanced Security in which the edauth -R command will not write user-profile entries to the root partition.
- Customers using C1crypt for password encryption on Enhanced Security systems have found that users are unable to change their passwords. Warnings from the passwd command state "Password not changed: failed to write protected password entry."

Number: Patch 1939.00**Abstract:** Feedback facility for dd**State:** New

- Provides the feedback facility for dd in long copies.
-

Number: Patch 1950.00

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, 1357.00, 1597.00, 1598.00, 1600.00, 1944.00, 1945.00, 1946.00, 1947.00, 1948.00

- 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 Tru64 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.
 - Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of passwords that have a length outside of the intended range.
 - 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 API calls when ASU is in use on the server.
 - Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file access.
 - Corrects a condition in which the useradd command was not managing default and template data properly. This shows up most notably with useradd -p producing the message "Password must be between 32 and 80 characters."
 - Corrects a problem in which extra leading/trailing spaces from the FIND dialog box in dxaccounts were not trimmed before being used, which resulted in a search failure.
 - Changes dxaccounts to cause it to display account expiration date at first view.
 - Corrects a condition in which the useradd error message is confusing when ASU is installed and the user runs useradd on a non-PDC server.
 - Corrects an userdel -D failure to remove a user from /etc/passwd on systems running Enhanced Security systems.
 - Includes the fix for the situation where dxaccounts does not allow to change the local user password in enhanced security environment.
 - Fixes a problem that prevented the locking and unlocking of multiple selected users under C2 security.
-

Patch 1950.00 Continued

- Corrects a problem in which running the usermod command with certain options such as -x distributed, grace_limit, gives a “Cross-device link” error when the var and yp directories belong to two different filesets on a NIS master.
- Prevents core dumps caused by the dxaccounts application in the following two cases:
 - When UNIX login name is modified and its corresponding PC account is not remapped to the new login name.
 - When a user tries to modify both UNIX account login name and PC account login name and tries to see information in ‘PC user view’.
- Corrects a problem in which dxaccounts core dumps when the root user is selected from the modification drop-down.

Number: Patch 1952.00**Abstract:** (SSRT0788U, SSRT0753U, SSRT0752U)**State:** Supersedes Patches 493.00, 495.00, 725.00, 1328.00

- Corrects a potential security vulnerability which, under certain circumstances, could compromise system integrity. This may be in the form of large values of ENVIRONMENT variables.
- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised when a buffer overflow occurs in the libXm utility. Buffer overflows are sometimes exploited in an attempt to subvert the function of a privileged program and possibly execute commands at the elevated privileges if the program file has the setuid privilege.
- Fixes a libXm.so incompatibility.
- Fixes a problem with the Motif ToggleButton Widget where, in some cases, it may not draw itself correctly.
- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file or privilege management.
- Corrects a problem in which the XmCvtXmStringToCT() function, under certain circumstances, does not correctly convert a compound string to a string in compound text format.

Number: Patch 1954.00**Abstract:** Security (SSRT0788U, SSRT0753U, SSRT0752U)**State:** Supersedes Patches 496.00, 498.00, 727.00, 1330.00

- Fixes a libXm.so incompatibility.
- Fixes a problem with the Motif ToggleButton Widget where, in some cases, it may not draw itself correctly.
- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised when a buffer overflow occurs in the libXm utility. Buffer overflows are sometimes exploited in an attempt to subvert the function of a privileged program and possibly execute commands at the elevated privileges if the program file has the setuid privilege.
- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file or privilege management.
- Corrects a problem in which the XmCvtXmStringToCT() function, under certain circumstances, does not correctly convert a compound string to a string in compound text format.

Number: Patch 1956.00**Abstract:** Resolution of a symlink problem**State:** New

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

Number: Patch 1958.00

Abstract: Correction to local and remote security domain risks

State: Supersedes Patches 1154.00, 1615.00

- Corrects 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.
- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file or privilege management.
- Corrects a potential security vulnerability that may result in unauthorized Privileged Access or a Denial of Service (DoS). This potential vulnerability may be in the form of local and remote security domain risks.

Number: Patch 1960.00

Abstract:awk not processing input files given in BEGIN section

State: Supersedes Patch 1671.00

- Fixes a situation, in which awk was not processing the input files specified in the BEGIN section. The awk/nawk command shows expected behavior with pipes.
- This patch fixes three issues for awk:
 - Corrects a problem of system integrity being compromised. This may be in the form of improper file or privilege management.
 - Returns the appropriate exit status when a disk is full.
 - Accepts input records of lengths up to 5,119 bytes.

Number: Patch 1962.00

Abstract: Corrects improper file access

State: Supersedes Patch 88.00

- Fixes a problem with the disklabel command. Disklabel was displaying large unsigned values as negative numbers.
- Corrects a potential compromise to system integrity. This may be in the form of improper file access.

Number: Patch 1964.00

Abstract: Correction to improper file access

State: New

- Corrects a potential compromise to system integrity. This may be in the form of improper file access.

Number: Patch 1966.00

Abstract: Security (SSRT0638U)

State: Supersedes Patches 139.00, 141.00

- Allows the dxsetacl utility to delete access ACLs.
 - Corrects 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.
 - Corrects a potential compromise to system integrity. This may be in the form of improper file access.
-

Number: Patch 1968.00

Abstract: Fixes a fatal error in the spike command

State: Supersedes Patches 150.00, 555.00, 1188.00, 1190.00, 803.00, 1658.00

- Fixes a problem that may cause the third command and other Atom-based instrumentation tools to fail.
- Fixes a fatal assertion error reported by pixie, hiprof, third spike, cord, uprofile and odump object file tools for some executables linked at optimization level 2 (-O2) or greater.
- This patch fixes the following problems in the linker (/usr/bin/ld):
 1. A problem with the datatype of the linker-defined _fpdata symbol.
 2. A problem that causes a linker crash when certain data alignment directives are used in the link.
 3. A problem in which the linker may corrupt the shared object registry file when -update_registry is specified with concurrent links.
 4. An error that occurs when the command ld -update_registry /dev/null is specified.
 5. A 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.
- Corrects a condition that causes a fatal error in /usr/bin/spike.

Number: Patch 1970.00

Abstract: Fixes a fatal error in the spike command

State: Supersedes Patches 158.00, 1198.00, 1200.00, 1660.00

- Fixes a problem where the spike command may fail to delete the low instruction of a pair of related instructions, causing it to abort with a runtime error.
- Corrects a problem in which prof -pixie -testcoverage <exe><exe>.Counts sometimes reports invalid source line number ranges.
- Fixes performance tool failures on Sierra Clusters (PFS) Parallel File Systems.
- Fixes a fatal assertion error reported by pixie, hiprof, third spike, cord, uprofile and odump object file tools for some executables linked at optimization level 2 (-O2) or greater.
- Corrects a condition that causes a fatal error in /usr/bin/spike.

Number: Patch 1972.00

Abstract: vmstat incorrect kernel per-RAD memory usage

State: New

- Corrects the output of the vmstat command for per-RAD kernel usage data.

Number: Patch 1974.00

Abstract: Updates the mcutil .h files

State: Supersedes Patch 1212.00

- Fixes and improves the mcutil program by correcting how bus resets are handled by the program, and by enhancing its error reporting capabilities.
- Distributes the latest .h files for the mcutil program.

Number: Patch 1976.00

Abstract: Fix for nvbmtpg -v option

State: New

- Enables the nvbmtpg utility to display only all the data in a specified mcell if the -v option has been selected.
-

Number: Patch 1978.00

Abstract: dxproctuner not displaying any data in fr_FR locale

State: Supersedes Patches 438.00, 1802.00

- 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.
- Corrects a problem with dxproctuner that caused it to give an error message or dump core when invoked.
- Corrects a condition in which dxproctuner does not display data in the fr_FR locale

Number: Patch 1980.00

Abstract: Local and remote security domain risks

State: New

- Corrects a potential security vulnerability problem that may result in unauthorized Privileged Access or a Denial of Service (DoS). This potential vulnerability may be in the form of local and remote security domain risks.

Number: Patch 1982.00

Abstract: Potential security vulnerability in gated

State: Supersedes Patches 1222.00, 1642.00, 1643.00, 1645.00, 1819.00

- Corrects a problem using MD5 authentication with Version 2 RIP
- Resolves a problem where the cluster interconnect route is inappropriately advertised.
- Corrects a problem in which gated, upon route aging timeout, deletes a static or loopback host route that is added via routing socket. In a cluster alias environment, this problem causes all TCP/UDP packets destined to the cluster alias address to mishandled.
- Adds support to gated for the option aliases-next-hop. This option provides a preference for the selection of next-hop address when multiple addresses exist on the interface.
- Resolves a problem with gated where adding a route may not succeed under certain circumstances.
- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may occur when Gated daemon or Gated Control utility (gdc) incorrectly accesses temporary files.

Number: Patch 1984.00

Abstract: Modification to secconfig suitlet

State: Supersedes Patch 749.00

- Makes the customize database option available when using secconfig for shadow passwords.
- Eliminates two warning messages from the SysMan Security Configuration tool that appear under the following conditions:
 - When changing the root password on systems with ASU installed.
 - When a user tries to trim auth logs before they are even created.

Number: Patch 1986.00

Abstract: Suppress an in_cksum() debug message sent to console

State: New

- Causes an in_cksum() console debug message to be suppressed.
-

Number: Patch 1988.00

Abstract: Fixes for the collect utility

State: Supersedes Patches 175.00, 799.00, 1238.00, 1240.00, 1654.00, 1656.00

- 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.
- Fixes the way collect handles Floating Point Exception.
- Fixes a problem in which collect was unable to create a new data file from a data file that does not include a termination record.
- Corrects a problem in which the collect utility exits with the message "chmod or chown failed.: No such file or directory" when it is run with the -H option and the directory linking to /var/adm/collect.dated does not exist.

Number: Patch 1990.00

Abstract: Fix for printer problems

State: New

- Fixes a problem in which the user cannot connect to a printer on the parallel port after cancelling a print job.

Number: Patch 1992.00

Abstract: Fix for fixfdmn problems

State: Supersedes Patches 177.00, 559.00, 818.00, 819.00, 820.00, 822.00, 1300.00, 1302.00, 1689.00, 1691.00

- 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.
 - Fixes three rare cases where fixfdmn could either fail to correct a domain or incorrectly make changes to a valid domain.
 - Allows fixfdmn to fix a rare corruption case in the RBMT/BMT0.
 - Improves the capability of fixfdmn to fix invalid extent data.

Number: Patch 1994.00

Abstract: Security (SSRT2280)

State: Supersedes Patches 561.00, 547.00, 813.00, 1284.00, 1285.00, 1287.00

- Fixes a problem with cut-and-paste operations of JISX0212 Japanese characters on X Window System applications.
- Fixes a problem in the X Toolkit library (Xt) that could cause the TeMIP Iconic_map Presentation Module application (mcc_iconic_map) to crash.
- Fixes a problem in which the definition of the X Toolkit function XtPending() was changed, thereby causing some applications built on earlier versions of Tru64 UNIX to fail.
- 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.
- Fixes a problem where a Chinese character whose byte sequence contains 0x9b cannot be entered with dxhanziim or cut-and-pasted.

Number: Patch 1996.00

Abstract: Security (SSRT2280)

State: Supersedes Patches 549.00, 815.00, 563.00, 1288.00, 1289.00, 1291.00

- Fixes a problem in the X Toolkit library (Xt) that could cause the TeMIP Iconic_map Presentation Module application (mcc_iconic_map) to crash.
- Fixes a problem in which the definition of the X Toolkit function XtPending() was changed, which caused some applications built on earlier versions of Tru64 UNIX to fail.
- Fixes a problem with cut-and-paste operations 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.
- Fixes a problem where a Chinese character whose byte sequence contains 0x9b cannot be entered with dxhanziim or cut-and-pasted.

Number: Patch 1998.00

Abstract: Modified volplex to detach mirrored swap plex

State: New

- Changes volplex to allow a mirrored swap plex to be detached when it contains a phantom subdisk.

Number: Patch 2000.00

Abstract: munlockall() locks process memory

State: Supersedes Patch 136.00

- Corrects the behavior of munlockall in the realtime library (librt).
-

Number: Patch 2002.00

Abstract: Static library fix for librt

State: Supersedes Patch 138.00

- Corrects a problem in which the behavior of POSIX 4 message queues does not follow the standard and returns unique message descriptors.
- Corrects the behavior of munlockall in the realtime library (librt).

Number: Patch 2004.00

Abstract: Fix incorrect diagnostic msg in traceroute

State: Supersedes Patch 718.00

- 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.
- Fixes an incorrect diagnostic message in the traceroute command.

Number: Patch 2007.00

Abstract: Assembler treats octal constants as decimal values

State: Supersedes Patches 471.00, 720.00, 2005.00

- Installs Version 3.06.08 of the assembler and resolves 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.
- When a load instruction specifies a target register and a base register that are the same, the assembler may generate an invalid optimization.
- Fixes a yacc stack overflow error in the Tru64 UNIX assembler.
- Provides version 3.06.10 of the assembler. This version fixes a problem encountered in version 3.06.09, wherein the assembler incorrectly treats octal constant data as if it were decimal.

Number: Patch 2009.00

Abstract: Potential security vulnerability in mrouted

State: New

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of the mrouted daemon incorrectly creating and accessing temporary files.

Number: Patch 2011.00

Abstract: lsmbstartup fix to ignore comments in /etc/fstab

State: New

- Fixes the situation in which lsmbstartup did not ignore comment lines in /etc/fstab. It was therefore possible for it to attempt to process that line if it appeared to be an entry for the root file system.

Number: Patch 2013.00

Abstract: Corrects improper file access

State: Existing

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

Number: Patch 2015.00

Abstract: Resolves a symlink problem

State: New

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

Number: Patch 2017.00

Abstract: Corrects improper file access

State: New

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

Number: Patch 2019.00

Abstract: Prevents garbled characters with dxdw

State: New

- Fixes a problem where, under certain circumstances, the Display Window application (dxdw) displays garbled characters in the transcript area of the application.

Number: Patch 2021.00

Abstract: Corrects improper file access

State: New

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

Number: Patch 2023.00

Abstract: Corrects improper file access

State: New

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

Number: Patch 2025.00

Abstract: Corrects improper file access

State: New

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

Number: Patch 2027.00

Abstract: Security (SSRT1-45U)

State: Supersedes Patch 1226.00

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

Number: Patch 2029.00

Abstract: Fix for Null Pointer Exception in LSMSA GUI

State: Supersedes Patch 455.00

- Addresses a problem in the display of disk controller to disk hierarchy by the lsmsa product.
- Fixes the problem in LSMSA GUI where the application throws a Null Pointer Exception while trying to display Volume/Fileset information on systems having ASU file-on-file mounts.

Number: Patch 2031.00

Abstract: Corrects improper file access

State: New

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

Number: Patch 2033.00

Abstract: Security (SSRT2193)

State: Supersedes Patch 1237.00

- Corrects 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.
- Corrects a potential security vulnerability that may result in unauthorized Privileged Access or a Denial of Service (DoS). This potential vulnerability may be in the form of local and remote security domain risks.

Number: Patch 2035.00

Abstract: Security (SSRT2280)

State: Supersedes Patches 1249.00, 1251.00

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

Number: Patch 2037.00

Abstract: send disk group option to usage utils for volplex

State: New

- Fixes a problem in which the volplex command was not sending the disk group option specified by the user to the usage type utility when both the usage type and disk group option are specified on the command line.

Number: Patch 2039.00

Abstract: Corrects improper file access

State: New

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

Number: Patch 2057.00

Abstract: Security (SSRT0788U, SSRT0753U, SSRT0752U)

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, 1008.00, 1560.00, 1561.00, 1563.00, 1934.00, 1935.00, 1937.00

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. The ttddbserverd 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 access.
- Fixes the problem of palette files not being read from /etc/dt/palettes.
- Fixes the dtprintinfo memory fault problem with long LANG value.
- Corrects 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.
- Corrects a potential security vulnerability in CDE Subprocess Control Service (dtspecd), which has a potential buffer overflow condition that may lead to unauthorized access.
- Fixes several potential security vulnerabilities where, under certain circumstances, system integrity may be compromised. This may be in the form of large values of command line arguments.
- Fixes a problem where a CDE session hangs at startup using localized .dt files located in the ~/.dt/types directory.
- Fixes several potential security vulnerabilities where, under certain circumstances, system integrity may be compromised. This may be in the form of improper privilege management.
- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised when a buffer overflow occurs in the DtSvc utility. Buffer overflows are sometimes exploited in an attempt to subvert the function of a privileged program and possibly execute commands at the elevated privileges if the program file has the setuid privilege.
- Corrects a potential security vulnerability in CDE that may result in unauthorized privileged access. The potential vulnerability may be in the form of local and remote security domain risks.

SSRT3589 - dtmailpr

Number: Patch 2061.00

Abstract:Correct handling of IPv6 literal addresses in rcp

State: New

- Corrects a problem with handling of IPv6 literal addresses in rcp.
-

Number: Patch 2064.00

Abstract: Modification to vdump and vrestore archive programs

State: Supersedes Patches 479.00, 119.00, 451.00, 3.00, 5.00, 447.00, 449.00, 113.00, 115.00, 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, 1040.00, 1575.00, 1578.00, 1579.00, 1580.00, 1581.00, 2062.00

- Corrects a problem in which access to a file may be denied when multiple processes attempt to access the same file at the same time, and access to the file should be allowed by an ACL on the file.
 - Corrects a problem in which if the ACL on a file is corrupted, the corrupted ACL is passed into the kernel causing a variety of problems.
 - Corrects a problem in which an AutoFS intercept point for a direct map entry may no longer induce automounts after an error has been detected during a previous automount attempt.
 - Fixes AutoFS problems as follows:
 - Eliminates error messages concerning property lists seen via certain utilities such as vdump.
 - Causes AutoFS automounts 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.
 - Prevents a core dump from vdump when a message length is greater than MAX_MSG_SIZE.
 - 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.
 - 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 hard-linked files is incomplete when they exist in different directories (fails to create directory for second occurrence of file with same inode number).
 - Eliminates inefficient behavior by autofs when the top-level directory of a direct hierarchical automount map entry cannot be successfully mounted.
 - Ensures that AutoFS correctly uses the mount options specified in automount map entries with replicated servers.
 - Fixes a problem where the tar -F (Fasttar) 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 in which the aufosd and autofs mount daemons do not properly parse the wildcard (*) character 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.
-

Patch 2064.00 Continued

- Fixes a one byte gap in the maximum size in the tar command before an extended header record is used (8589934591 (octal 7777777777)).
 - Eliminates AutoFS automount 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 automounted file systems may be mounted more than once.
 - Ensures that autofs mount will process all NIS-provided map files included via the plus sign (+) 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 automount map entry.
 - Allows AutoFS to correctly handle NIS-based automount 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 automount failures.
 - Changes vrestore as follows:
 - A file or directory name is displayed along with the error message when command fails to set a property list
 - It will not dump core when a tape has a smaller blocksize than expected.
 - It handles no-rewind tapes properly.
 - It reads environment variable for user-defined device name.
 - It allows attributes to be set to the top level directory.
 - Fixes tar to properly handle unusual directory specifications.
 - Corrects the tar program to properly handle unusual directory specifications.
 - Makes the following changes to vdump and vrestore:
 - Fixes vdump to receive a Ctrl/c interrupt in the expected way.
 - Fixes the vdump to pick up correct messages in all locales.
 - Causes vdump to avoid some unnecessary function calls, thus allowing faster vdumps.
 - Fixes vrestore to receive a Ctrl/c interrupt in the expected way.
 - Fixes vrestore to pick up correct messages in all locales.
 - Fixes vrestore to display bit file attributes upon -l option.
 - Prevents vrestore from failing during a remote system call.
 - Increases the maximum blocksize for vdump and vrestore.
 - Fixes a problem that could cause vdump to crash due to compression code overflow.
 - Disallow block sizes greater than 64 K. This reverses a feature that was added in Patch Kit 5 that caused forward compatibility problems.
 - Allows vdump to properly backup less than 512 byte files from CD-ROM.
 - Causes the programs to avoid a cluster drd bandwidth problem that occurs when a restore is attempted via a private tape drive that is not on the current node.
 - Causes the programs to avoid misinterpretation of archive version, which would happen in certain locales, causing vrestore to abort.
 - Reports errors emerging from the failure of the close() call for devices/files which are on an NFS mount point.
-

Number: Patch 2066.00

Abstract: Revises the sys_attrs_inet.5 reference

State: New (Supersedes Patch 1705.00)

- Revises the sys_attrs_inet.5 reference page to document the change of the increased default values for udp_ttl and tcp_ttl to 128 hops. and revises the sys_attrs_net(5), netstat(1), and route(8) reference pages.
- Revises the sys_attrs_net(5), netstat(1), and route(8) reference pages.

Number: Patch 2068.00

Abstract: Resolves a problem that causes rsh process hangs

State: New

- Resolves a problem that can cause the rsh processes to hang.

Number: Patch 2070.00

Abstract: Corrects improper file or privilege management

State: Supersedes Patches 1722.00, 1723.00, 1725.00

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

Number: Patch 2072.00

Abstract: Revises volrestore(8) and volsave(8) reference pages

State: New

- Revises the volrestore(8) and volsave(8) reference pages to document a change that allows volsave and volrestore to save nconfig and nlog policies for diskgroups and restore them appropriately.

Number: Patch 2078.00

Abstract: Security fix for sendmail utility

State: New (Supersedes Patches 1662.00, 1777.00)

- Corrects a potential security vulnerability in sendmail which may result in non-privileged users gaining unauthorized access to files or privileged access on the system. This potential vulnerability may be in the form of a local or remote security domain risk.
- Corrects several potential security vulnerabilities that may result in unauthorized Privileged Access or a Denial of Service (DoS). These vulnerabilities may be in the form of local and remote security domain risks.

SSRT3631 sendmail - (Severity - High)

Number: Patch 2084.00

Abstract: Security (SSRT2301, SSRT0845U, SSRT2266, SSRT2322)

State: Supersedes Patches 160.00, 1065.00, 1159.00, 1160.00, 1161.00, 1162.00, 1164.00, 1352.00, 1264.00, 222.00, 67.00, 582.00, 574.00, 576.00, 703.00, 839.00, 1012.00, 1013.00, 1014.00, 1015.00, 1017.00, 1344.00, 1346.00, 269.00, 1025.00, 232.00, 504.00, 1107.00, 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, 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, 1367.00, 125.00, 309.00, 460.00, 461.00, 463.00, 734.00, 735.00, 737.00, 1130.00, 1131.00, 1132.00, 1134.00, 97.00, 84.00, 1170.00, 1379.00, 1380.00, 1381.00, 1382.00, 1383.00, 1384.00, 1385.00, 1386.00, 1387.00, 1388.00, 1389.00, 1390.00, 1391.00, 1392.00, 1393.00, 1394.00, 1395.00, 1396.00, 1397.00, 1398.00, 1399.00, 1400.00, 1401.00, 1402.00, 1403.00, 1404.00, 1405.00, 1406.00, 1407.00, 1408.00, 1409.00, 1410.00, 1411.00, 1412.00, 1413.00, 1414.00, 1415.00, 1416.00, 1417.00, 1418.00, 1419.00, 1420.00, 1421.00, 1422.00, 1423.00, 1424.00, 1425.00, 1426.00, 1427.00, 1428.00, 1429.00, 1430.00, 1431.00, 1432.00, 1433.00, 1434.00, 1435.00, 1436.00, 1437.00, 1438.00, 1439.00, 1440.00, 1441.00, 1442.00, 1443.00, 1444.00, 1445.00, 1446.00, 1447.00, 1448.00, 1449.00, 1450.00, 1451.00, 1452.00, 1453.00, 1454.00, 1455.00, 1456.00, 1457.00, 1458.00, 1459.00, 1460.00, 1461.00, 1462.00, 1463.00, 1464.00, 1465.00, 1466.00, 1467.00, 1468.00, 1469.00, 1470.00, 1471.00, 1472.00, 1473.00, 1474.00, 1475.00, 1476.00, 1477.00, 1478.00, 1479.00, 1480.00, 1481.00, 1482.00, 1483.00, 1484.00, 1485.00, 1486.00, 1487.00, 1488.00, 1489.00, 1490.00, 1491.00, 1492.00, 1493.00, 1494.00, 1495.00, 1496.00, 1497.00, 1498.00, 1499.00, 1500.00, 1501.00, 1502.00, 1503.00, 1504.00, 1505.00, 1506.00, 1507.00, 1508.00, 1509.00, 1510.00, 1511.00, 1512.00,

Patch 2084.00 Continued

1513.00, 1514.00, 1515.00, 1516.00, 1517.00, 1518.00, 1519.00, 1520.00, 1521.00, 1522.00, 1523.00, 1524.00, 1525.00, 1526.00, 1527.00, 1528.00, 1529.00, 1530.00, 1531.00, 1532.00, 1533.00, 1534.00, 1535.00, 1537.00, 1602.00, 1781.00, 1787.00, 1788.00, 1789.00, 1790.00, 1791.00, 1792.00, 1793.00, 1794.00, 1796.00, 1805.00, 1806.00, 1807.00, 1808.00, 1809.00, 106.00, 1179.00, 1727.00, 1811.00, 1826.00, 1827.00, 1828.00, 1830.00, 195.00, 1315.00, 1832.00, 1833.00, 1834.00, 1835.00, 1837.00, 1838.00, 1839.00, 1840.00, 1841.00, 1842.00, 1843.00, 1844.00, 1845.00, 1846.00, 1847.00, 1848.00, 1849.00, 1850.00, 1851.00, 1852.00, 1853.00, 1854.00, 1855.00, 1856.00, 1857.00, 1858.00, 1859.00, 1860.00, 1861.00, 1862.00, 1863.00, 1864.00, 1865.00, 1866.00, 1867.00, 1868.00, 1869.00, 1870.00, 1871.00, 1872.00, 1873.00, 1874.00, 1875.00, 1876.00, 1877.00, 1878.00, 1879.00, 1880.00, 1881.00, 1882.00, 1883.00, 1884.00, 1885.00, 1886.00, 1887.00, 1888.00, 1889.00, 1890.00, 1891.00, 1892.00, 1893.00, 1894.00, 1895.00, 1896.00, 1897.00, 1898.00, 1899.00, 1900.00, 1901.00, 1902.00, 1903.00, 1904.00, 1905.00, 1906.00, 1907.00, 1908.00, 1909.00, 1910.00, 1911.00, 1912.00, 1913.00, 1915.00, 2040.00, 2041.00, 2042.00, 2043.00, 2044.00, 2045.00, 2046.00, 2047.00, 2048.00, 2049.00, 2050.00, 2051.00, 2052.00, 2053.00, 2055.00, 2074.00, 2076.00, 2079.00, 2080.00, 2081.00, 2082.00

- 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.
 - Adds support for new 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.
 - Corrects a potential security vulnerability which, under certain circumstances, could compromise system integrity. This may be in the form of improper file access.
 - 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.
 - 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 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.
-

Patch 2084.00 Continued

- 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.
- Fixes a volrecover error of "Cannot refetch volume" when volumes exist only in a non-rootdg diskgroup.
- Moves 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.
- Corrects 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)

- Fixes a problem in NetRAIN by causing NetRAIN interface creation to fail if any of the requested standby interfaces do not exist.
 - Adds support to ifconfig application for the IPv6 command line argument ip6reachabletime.
 - 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.
 - 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).
-

Patch 2084.00 Continued

- 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 failed due to the PARSING of ps command output.
 - Fixes a multi-thread timing window in malloc and free where the list of free chunks could become corrupted, resulting in a segfault.
 - Fixes a regular expression performance problem in sed.
 - Fixes a problem with printing long double values.
 - Fixes a performance degradation in malloc, in applications which perform many mallocs and few frees. With this patch, the performance of malloc is constant regardless of the number of allocated chunks outstanding.
 - Fixes a problem with atexit() or pthread_atfork() handlers in shared libraries. An application will crash when handlers in shared libraries are called after the libraries are 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.
-

Patch 2084.00 Continued

- 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 panicking 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.
 - 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.
-

Patch 2084.00 Continued

- Fixes a performance problem and the results are large performance increases in configurations where more than eight tapes are supported on a FibreChannel (usually behind an MDR or FCTCII).
 - 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 (8 MB or larger) System V Shared memory regions are in use.
 - 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 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).
-

Patch 2084.00 Continued

- 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 the 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.
 - 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).
 - 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().
 - 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.
 - Provides AdvFS and VFS support for the freezefs and thawfs commands.
 - Provides EVM V2 enablers.
 - Enables access to SCSI control ports (/dev/cport/scp??), allowing management of some types of RAID controllers.
 - Eliminates unintended AutoFS automount 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.
-

Patch 2084.00 Continued

- 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.
 - Addresses three UBC issues:
 - Reinstates ubc_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.
-

Patch 2084.00 Continued

- 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.
 - 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.
-

Patch 2084.00 Continued

- 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.
- 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
```

Patch 2084.00 Continued

- 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-keygen.
 - 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 node 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 mount points in the cluster.
 - Eliminates unintended AutoFS automounts, 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_XNTS.
 - 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.
 - 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 behavior 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.
-

Patch 2084.00 Continued

- 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, there is a domain panic on reboot.
 - 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 the LUN is greater than 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.
 - 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 FibreChannel 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 2084.00 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 file system 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 2084.00 Continued

- Fixes an ISO9660 file system size limitation of 2.1 GB 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 that results in broadcast or multicast packets being processed multiple times on behalf of a NetRAIN device, once for each backup interface.
 - Fixes a problem that causes the 4.3BSD socket interface to return incorrect values for IOCTL calls accessing IP alias address information.
-

Patch 2084.00 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 system with more than one CPU.
 - 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 UNIX 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 2084.00 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) hangs 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 8 K.
 - 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 2084.00 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 preemption 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 2084.00 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 discontinuous 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 2084.00 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 performing a copy-on-write operation 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 (particularly 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 2084.00 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 UNIX 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 identified 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 occurring 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 CD-ROM 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 getusage function where the memory usage could be incorrectly reported.
 - Fixes a problem of an occasional deadlock during process exit for multithreaded processes.
-

Patch 2084.00 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 2084.00 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.
 - 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.
-

Patch 2084.00 Continued

- 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 Version 6 (NHD-6) devices.
 - 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.
 - 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 that occurs when a ksh menu is started from within user's .profile file and ksh does 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.
 - Fixes a problem where the vi editor core dumps when it finds invalid syntax during a substitute operation.
 - Fixes a problem in which shutting down the network would also shut down the cluster interconnect interface in a LAN cluster.
 - Corrects several potential security vulnerabilities which, under certain circumstances, could compromise system integrity. These may be in the form of improper file or privilege management.
 - Prevents vold from core dumping when attempting to delete a disk that was not initialized properly.
 - Fixes a deadlocking problem in the kernel where a file open on a clone could hang when ACLs are enabled.
 - Fixes a problem where gh_min_seg_size could not be set below 8 M.
 - Corrects a KMF caused by uninitialized or incorrect parameters passed to the setsockopt system calls.
-

Patch 2084.00 Continued

- Fixes a problem where Smart Array 5300 Logical Volumes were counted as RAID controllers.
 - Installs Version 1.12 of the ciss driver.
Previous releases of the Version 5.1A patch kit installed versions 1.02, 1.07, 1.08, 1.09, and 1.11 of the ciss driver.
 - Fixes a problem that can cause a GS80 to hang after a SA5300 reset.
 - Fixes a potential system crash when shutting down after using a DAPBA or DAPCA ATM adapter.
 - Corrects a potential problem on a system generating a heavy volume of audit events in which the auditd -d command that flushes the kernel audit buffers could cause audit data corruption on a multi-CPU machine.
 - Fixes a problem that can cause a hang to occur during the renaming of an AdvFS file.
 - Prevents a kernel memory fault panic in _OtsMove when going through the fs_read() routine.
 - Fixes a casting error in AdvFS that produces a warning message.
 - Fixes potential crashes from within the pshared subsystem.
 - Fixes the changer driver to report the manufacturer ID.
 - Fixes nissetup to leave /etc/group with an incorrect mode of 600 after removing NIS.
 - Corrects a newfs and extendfs problem in which a statically sized cylinder summary area of a UFS file system limited the overall size of the file system. In some situations in which a large (greater than 800 GB) file system is created, newfs would not exit with an error even if the cylinder summary area could not be written to the disk.
 - Corrects a problem with mfs in which requests to create a memory file system with a large size can result in a core dump and silent error.
 - Adds system configuration tunables for AlphaServer ES45 environmental monitoring.
 - Fixes a problem in file property lists where the maximum length of a property list name was limited to 245 characters. The new limit is 255.
 - Reschedules certain default cleanup cron jobs so that they are not skipped during a time change to DST.
 - Prevents a possible panic in bs_derefpq.
 - Change an rws write lock in the VMAC lookup routine to an rws read lock for better SMP scaling.
 - Fixes a regression from pre-Version 5.0 releases in the libc mktime() function's handling of potentially ambiguous tm struct times; that is, those that fall within a backward clock shift and have an initially negative tm_isdst value).
 - Fixes a cluster panic with the following error message:

```
panic (cpu 3): ics_unable_to_make_progress: heartbeat checking blocked
```
 - Fixes a system panic that produces one of the following panic strings:

```
mcs_lock: lock already owned by cpu  
thread_block: simple lock owned
```
 - Fixes a problem encountered with the Bourne shell when a file name with trailing slash (/) is used as an argument to the command.
 - Corrects problems in tape open and tape ioctl when the FNDELAY flag is set, which could result in tape devices not responding or in a kernel memory fault.
 - Fixes a condition that results in system hangs or panics resulting from bad locking in the AdvFS msfs_getpage routine.
 - Fixes a problem where the home directory and login shell attributes for a user account were not supplied to the audit daemon for authentication failures.
-

Patch 2084.00 Continued

- Adds recognition for possible future devices.
- Eliminates compiler warnings in ksh.
- Increases the TCP window from 96 KB to 500 KB to improve performance.
- Changes the netisr thread to dynamically estimate the reply size and reserve the space in the socket buffer.
- Adds a new timeout check to notice when the data has not been acknowledged in 30-50 seconds and copy those buffers. This allows the UBC to free up those mbufs.
- Corrects a problem when running Enhanced Security in which an emergency log in for the root account on the console would fail in TruCluster configurations with the following string:

Impossible to execute /sbin/sh

- Provides a workaround for a domain panic when corruption in the deferred-delete list of an AdvFS file system is detected.
 - Changes the /usr/sbin/dirclean utility so it no longer attempts to remove the AdvFS .tags directory or the quota.group and quota.user files.
 - Prevents kmfs when AdvFS encounters proplist metadata inconsistencies.
 - Fixes a system panic in the ubc_page_stealer routine.
 - Corrects a kernel memory fault in the table syscall.
 - Improves kernel module functions
 - Fixes various problems in the libc functions getdate(), strptime(), callrpc(), strncasecmp(), fork(), and popen(), and the libnuma function ncreate().
 - Fixes a cross-node cluster deadlock that can occur when AdvFS threads on two cluster nodes simultaneously call code that requires the taking of already held AdvFS locks on the other node.
 - Fixes a problem in which volmigrate returns a shell error when attempting to migrate an AdvFS domain with multiple filesets. These domains can now be migrated as long as all the filesets are mounted.
 - Fixes sh to print the correct msg when enhanced core file naming is on.
 - Enhances cron to do extensive logging.
 - Addresses an issue in which I/O may not be completed under certain circumstances.
 - Fixes client (login, su, rshd, edauth, and sshd2) hangs and long delays under Enhanced Security, as well as some intermittent errors or failures seen with prpasswd or rpc.yppasswdd.
 - Enables mountd to correctly handle entries with multiple lines input in exports file.
-

Patch 2084.00 Continued

- Fixes a problem in the alt driver that prevents DEGPA from being used with DE50x or DE60x adapters in a LAG set.
- Makes start up scripts in /sbin/init.d world readable.
- Provides protection against a class of potential security vulnerabilities by allowing a system administrator to enable memory management protections that limit potential buffer overflow vulnerabilities. Buffer overflows are sometimes exploited in an attempt to subvert the function of a privileged program and possibly execute commands at elevated privileges if the program file has the setuid privilege.
- Improves AdvFS scalability. These improvements help reduce lock contention and improve performance.
- Corrects lockmode 4 problem in cdisk_event_notify.
- Prevents a hang in LSM/CLSM commands in sosleep().
- Fixes a problem where the system can panic with a "kernel memory fault" in simple_lock(), being called from fuser().
- Fixes a possible (though rare) hang in the hardware configuration subsystem.
- Prevents a situation where an AdvFS Migrate syscall launched from the CFS client could result in a hung client thread and a hung server thread.
- Fixes a panic resulting from race condition in the memory file system (MFS) over the cluster file system (CFS).
- Fixes a problem in which rmvol causes file and directory information to become lost, which can cause domain panics or strange behavior on file statistics.
- Fixes a case where rmvol causes files to lose their ACLs, which opens up a security hole on those files.
- Fixes a case where rmvol enters into an infinite loop while trying to move a file from one volume to another.
- Fixes a ksh problem related to cleaning the process when terminal is abruptly stopped.
- Fixes a problem that could cause the following message to be incorrectly generated:

ccfg_MakeDeviceIdentWWID: Invalid device ID

- Fixes a condition that can cause the following erroneous console warning message to be provided when cluster root is under LSM control:

WARNING: cluster root devices are on private buses!

- Fixes a problem which causes a hang in fcntl()/setflock().
 - Prevents a kernel memory fault when an original fileset is unmounted during the deletion of its clone fileset.
 - Fixes an occasional panic that can be seen when reading from a process using Granularity Hints via procs.
 - Increases the default values for udp_ttl and tcp_ttl to 128 hops.
 - Changes the fwupgrade command to allow the specified firmware update image to be located on a BOOTP server in a connected network.
-

Patch 2084.00 Continued

- Fixes the following problems in sh:
 - A service denial problem when a quoted script is executed.
 - A problem with handling ELF files.
 - A problem with the shell variable \$- not holding -C set option when it is turned on.
 - A problem printing broken characters when the type builtin utility of sh is invoked in a Japanese locale.
 - Fixes the bcmdriver for DEGXA to correct problems that cause it to incorrectly report data overruns and prevent DEG2-SA modules from being recognized.
 - Fixes the IDE/ATAPI driver's reset logic to prevent a kernel memory fault when booting and to properly detect and log all master and slave reset failures when the system is operational.
 - Fixes the error "invalid IPL 4."
 - Fixes a problem in which the return value of an unlink() call is not checked when two threads try to move a file to two different destination, and although one of the threads could unlink() the source file, no relevant error message is displayed.
 - Improves performance for removing or truncating large files on UFS file systems.
 - Corrects an issue of a cluster node hanging on boot while the other member recovers the cluster root file systems.
 - Corrects a problem in which multi-CPU systems will sometimes live lock while processing incoming network traffic. In some cases the live lock could result in a cluster event timeout panic.
 - Addresses an issue where NULL inquiry data causes a "Device has no name" error as well as possible I/O stalls.
 - Prevents unnecessary retries of the following: On HSG80, fail unit attention with ascq = 0xf002 and returns proper error to a higher layer.
 - Corrects the improper scheduling of cron jobs related to months not having 31 days.
 - Fixes a problem where a device file, such as /dev/console, can become inaccessible, returning the error "Bad File Number."
 - Fixes a problem where attempts by the runtime loader (/sbin/loader) to free a null pointer are in error.
 - Allows volsave and volrestore to save nconfig/nlog policies for diskgroups and restore them appropriately.
 - Corrects awk errors for invalid quit statements.
 - Changes ICMP redirect processing code so it does not create an invalid route for an invalid redirect.
 - Prevents crontab from removing its entries and vi from truncating the existing file when a file system is full.
 - Fixes a problem where defragment (or migrate, balance, or rmvol) threads can hang in ms_malloc() and overlay_xtnt_map().
 - Fixes a problem in which NIS clients may fail to connect to non-Tru64 UNIX NIS servers that only support the V2 NIS protocol.
 - Fixes a problem where the CAM I/O subsystem does not always zero the Cam Control Blocks which are used by the peripheral drivers, which can cause a kernel memory fault or a system hang when the subsystem is low on memory.
 - Allows AdvFS to record if a domain panic has occurred, even if a system panic results.
 - Corrects a problem in which volsave would indicate that the configuration had changed during saving when nothing else was going on. This change allows for synchronization for a couple of vold requests and will make sure volsave and other commands are more cluster safe and synchronized.
-

Patch 2084.00 Continued

- Fixes a situation where, on a panic, the operating system will erroneously reboot instead of halt and fail to take a crash dump.
 - Corrects small memory leaks within the kernel that occur infrequently.
 - Eliminates the compiler warnings in ksh.
 - Adds a -M command option to newfs to let users specify permissions of an mfs root directory when it is first created.
 - Fixes a problem caused when the TCP layer prematurely closes a slow, but good connection with TCP reset.
 - Corrects an I/O performance issue seen when CLSM is configured and one member of a cluster goes down unexpectedly.
 - Fixes a potential kernel memory fault in the IPv6 subsystem. The problem is caused by ICMP error reporting in the IPv6 subsystems for packets containing IPv6 extension headers and options.
 - Fixes an sh problem while executing command substitution.
 - Corrects a locking problem with NFS running over UFS.
 - Verifies path structures in ctape_ioctl and ctape_generic_passthru to prevent a kernel memory fault if the tape was opened with FNDELAY flag set.
 - Prevents a panic that may occur in a cluster when a fileset mounted -o dual is failed over or unmounted during shutdown.
 - Corrects the creation of console boot device strings for devices on subordinate buses.
 - Fixes a problem in which an invalid character sometimes appears when requesting the name of a boot device via consvar.
 - Resolves internet protocol conformance issues and fixes a problem with sending multicast datagrams.
 - Resolves a problem where some de50x network interface cards, under specific circumstances, may not send gratuitous arp packets.
 - Clarifies the "LIDs do not match" error message. It now displays the values that do not match. The intent is to assist system managers or service personnel with troubleshooting when this error occurs.
 - Fixes library dependencies to allow ifconfig to be run in single user mode.
 - Fixes a slow boot when booting several cluster nodes at the same time and CLSM is configured.
 - Corrects a problem in which incorrect I/O status may be returned by the KZPEA driver when attempting to abort an I/O during a reset.
 - Corrects problems with the time of year (TOY) clock.
 - Adds support for IPV6_UNICAST_HOPS socket option on raw sockets.
 - Improve the fragment gathering mechanism to boost performance.
 - Fixes an AdvFS path which can cause a panic in the advfs_page_busy() routine.
 - Fixes a problem where memory could retain execute permission on EV6 machines.
-

Patch 2084.00 Continued

- Fixes a condition that causes a delete_pv_entry panic when kernel virtual address space has high usage.
- Fixes the /sbin/init.d/route script to correct a problem that causes routes to not get flushed properly.
- Corrects a potential security vulnerability that could result in a denial of service on RPC-based servers with Enhanced Security (C2) enabled. This vulnerability may be in the form of local and remote security domain risks.

SSRT2412 portmapper with Enhanced Security enabled (Severity - High)

- Corrects problems in UFS extendfs that could cause file system metadata inconsistency.
 - Corrects a potential security vulnerability that could result in nonprivileged users gaining unauthorized access to files or privileged access on the system. This potential vulnerability may be in the form of a local and remote security domain risk.
 - Fixes a problem where a process waiting on a semaphore does not get woken up.
 - Fixes a problem in which the quot -v command sometimes returns wrong quota information on UFS partition.
 - Fixes memory leaks caused by certain type of scripts.
 - Prevent a_lock related panics.
 - Fixes a problem where an I/O error (EIO) can occasionally be returned after a page fault.
 - Corrects netstat and ifconfig so that MAC addresses are printed using 2-digit hex octets with leading zeros.
 - Addresses system problems that can occur when the system is under heavy I/O load and/or low memory conditions.
 - Corrects a problem in which the class scheduler fails to restart if a system administrator starts then stops it and then removes system owned semaphores (via ipcrm -s). The error returned is "class_open: allocate or access semaphore Invalid argument."
 - Fixes a problem where a cluster member can panic with a kernel memory fault in kdm_isenabled().
 - Prevents a kernel memory fault panic that would occur when the audit daemon is set to periodically dump the kernel audit buffers to the audit log file (auditd -d freq).
 - Updates the camreport program to generate additional status and adds support for new devices.
-

Patch 2084.00 Continued

- Fixes an IDE/ATA bus hang caused by attempting to complete raw odd byte DMA transfers to/from IDE/ATAPI devices.
- Fixes a performance problem where threads could spend a long time in the check_busy() AdvFS routine.
- Increases the default limit of DLI packets to 16 K and makes the limit tunable.
- Modifies volencap to prevent encapsulation of restricted partitions or placing swap into non-rootdg diskgroup.
- Fixes a problem that could cause corruption in a forked process that may be evident when using the C shell.
- Provides protection against a class of potential security vulnerabilities called buffer overflows. Buffer overflows are sometimes exploited in an attempt to subvert the function of a privileged program and possibly execute commands at the elevated privileges if the program file has the setuid privilege. This patch allows a system administrator to enable memory management protections that limit potential buffer overflow vulnerabilities.
- ARP request for a permanent ARP entry is ignored, user cannot connect from remote system.
- Corrects a potential security vulnerability that could result in unauthorized privileged access or a denial of service. This potential vulnerability may be in the form of local and remote security domain risks.
- Corrects the following potential security vulnerability:

SSRT2384 rpc (Severity - High)

- Improves the scalability and performance of AdvFS.
 - Corrects a problem that causes an NFS client panic when it receives a null entry as a response to a readdirplus request from an NFS server.
 - Corrects a problem that causes an NFS client panic caused by receiving illegal file access mode from an NFS client.
 - Corrects a "blkfree: freeing free block" panic and a "blkfree: freeing free frag" panic.
 - Corrects a race condition that may result in hung disks under certain circumstances, for example, after a SCSI reset is done.
 - Prevents an IDE bus hang caused when issuing a play audio track command from scu to an ATAPI CD-ROM containing an enhanced CD.
 - Fixes an internal problem in the kernel's AdvFS, UFS, and NFS file systems where extended attributes with names greater than 247 characters could not be set on files. The new limit is 254 plus a null string terminator.
-

Patch 2084.00 Continued

- Corrects interoperability problems with the libsshrcmd.so library or any threaded library that uses the rcmd function.
 - Fixes a standard violation on AdvFS.
 - Replaces two potential panics in AdvFS with domain panics.
 - Corrects mv to correctly parse the pathname when source directory ends in trailing slash (/).
 - Prevents a potential hang during umount if a domain_panic has been encountered.
 - Fixes a problem in audit_tool which appends nonsense characters to the audit information to the output of an execve event in brief mode.
 - Prevents segmentation faults when sia_ses_init is passed a malformed argument vector.
 - Fixes the Bourne shell so that the expansion \$@ generates zero fields when no positional parameters are specified for the shell function.
 - Corrects a situation in which files created under directories such as /usr/lib/sabt/ during the btcreate of file systems with LSM gets copied on to the tape. The files will be removed before they are dumped to the tape.
 - Fixes a problem that causes the following panic:
panic: rdg: unwiring
 - Corrects a problem where, in some circumstances, a system may panic with a kernel memory fault when a device that is being opened by one program is being deleted via the hwmgr utility.
 - Fixes a potential panic in the auditing of the swapctl syscall.
 - Fixes a write/pwrite problem when O_SYNC is set.
 - Corrects a possible security hole reported by SSRT2323.
 - Allows the size of the NFS server's duplicate request cache to be adjusted as needed.
 - Corrects a problem in which logins in TruCluster environments using Enhanced Security could hang on any member other than the one serving /var to CFS.
 - Corrects a problem that could cause a cluster crash when VMAC is enabled.
 - Corrects memory file system size to accommodate the required space while creating bttape in AdvFS or UFS with LSM support.
 - Corrects a problem in which KZPEA firmware fails to correctly handle file marks with odd byte transfers.
 - Fixes a problem in the Network startup script that could result in a failure to configure an interface with an IP address.
 - Fixes an error in some sections of code that get an E_PAGE_NOT_MAPPED error when the code expected the page to exist.
 - Fixes two problems in the dynamic runtime loader that might cause an application to crash.
 - Closes a small race condition when accessing an internal data structure in AdvFS.
 - Causes restore to export file (-root=hostlist) behavior.
 - Fixes an error that can cause a multivolume domain to report ENO_MORE_BLKs when some volumes still have free storage.
 - Fixes an underlying problem in the NFS client that could lead to a panic on a single system or an assertion failure panic on a cluster.
-

Patch 2084.00 Continued

- Corrects the behavior of the sort command which now checks for duplicates with the -c, -u, and -k flags.
 - Changes sort to give exit value > 1 for all the error messages, in compliance with existing specifications.
 - Corrects the sort command so it does not dump core when more than 50 sort keys are used.
 - Fixes various problems in the ee driver for DE60x Ethernet adapters.
 - Fixes various problems in the bcm driver for DEGXA Gigabit Ethernet adapters.
 - Adds preliminary ddr support for the DAT-72 device.
 - Prevents kernel memory fault Panic from defragment or rmvol
 - Corrects a problem that could result in a system panic on close() if the BPF default packet filter is in use.
 - Corrects 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.
 - Corrects problems with name resolution when an error is encountered during processing of the local host files.
 - Allows tasks with just one SCS thread to be migrated in the same manner as single threaded processes in NUMA environments.
 - Improves null partition checking.
 - Corrects problems in AdvFS write logic, including:
 - a mismatch between the value reported by the write syscall and the number of bytes.
 - a problem that appears as unavailable, unused storage.
 - Provides the logging of an informational event to the errorlog when the peripheral driver is unable to get an user specified IDENTIFIER for a storage array device.
 - Modifies the ifconfig command to include an RCS ID for patch tracking purposes.
 - Fixes CAM errors that arise when opening CD-ROM devices containing blank media.
 - Fixes a problem where a cluster node hangs or panics during boot with a SLTO on the vdiLock. POSIX named semaphores will be in use on the cluster.
 - Corrects a problem in which Tru64 UNIX sees an HP-XP RAID array controller as a disk after an HP-XP storage device is added to the system.
 - Fixes the number of cylinders defined in the last cylinder group of a UFS file system that has been extended.
 - Corrects idle-wait time accounting within the AdvFS file system, as reported by the vmstat -w command.
 - Resolves a problem affecting IPv6 sockets where a packet transmission might fail when the destination address is not specified.
 - Changes the at command to prevent jobs that fail to execute when LANG and LD_LIBRARY_PATH are set.
 - Increases the auditability of memory wire and unwire operations.
-

Patch 2084.00 Continued

- Corrects a potential security vulnerability has been discovered where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file access.
 - Corrects a condition in which command response deteriorates to about 30 seconds when doing an rmvol on a domain with a volume containing large files (8 GB). The condition could last for the duration of the transfer, which could be 45 minutes.
 - Corrects a potential security vulnerability that may result in a local Denial of Service (DoS).
 - Fixes a problem in which data from an AdvFS file with a frag could be written to an incorrect location if an NFS client grew the file.
 - Provides comment handling capability in the route initialization script for parsing /etc/routes file.
 - Modifies the mountd daemon to address a core dump problem.
 - Updates the DDR entry for the MSA array and adds other future new device support.
 - Replaces several system panics with domain panics.
 - Enhances HP-XP array controller support and possible future new tape device support.
 - Fixes a problem in which clua.mod does not handle TCP RST messages appropriately.
 - Fixes a problem that can cause a kernel memory fault during boot if a FibreChannel LSM boot disk is not discovered during the device probe.
 - Allows acceleration of boot on large memory GS systems.
 - Provides modifications for possible future emx hardware.
 - Updates ICMP code to add redirect timeout support and nonmodifiable route table entries.
 - Provides hardware support for the DS15/TS15 platform.
 - Corrects a problem in which users are unable to change their passwords on systems running Enhanced Security that use C1crypt for password encryption. In such cases, the passwd command returns the warning "Password not changed: failed to write protected password entry."
 - Corrects a problem on systems running Enhanced Security in which the edauth -R command will not write user-profile entries to the root partition.
 - Fixes a "simple lock: time limit exceeded" panic.
 - Corrects the problem of the statfs() function returning EINVAL when operating on an fattach() function clone streams device.
 - Sets a software limit to prevent serious performance problems that can occur when TCP connections that have a rate-limit enforced by a downstream network device overrun the device.
 - Fixes a condition that could cause the panic "trap: invalid memory read access from kernel mode in procfs_psinfo()."
 - Fixes issues with I/O failures seen on non-rootdg disk groups, including the following:
 - The reimport of the disk group may fail under certain conditions
 - vold may core dump when failing a non-rootdg disk
 - Fixes a performance issue seen during plex detach operations.
-

Patch 2084.00 Continued

- Corrects a problem in which table() calls would not correctly get process arguments.
 - Corrects a panic resulting from a race condition between vnode deallocation logic and the use of CACHE_LOOKUP_REF/CACHE_LOOKUP_RELE in grab_bsacc and bs_dealloc_access (which attempt to block vnode deallocation).
 - Adds support for recognizing ELF and HP-UX file format to the file command.
 - Corrects a "simple_lock: time limit exceeded" issue that can occur in the CAM I/O subsystem when MCS locking is disabled.
 - Corrects the NFS server's handling of files open for direct I/O.
 - Fixes two problems in the dynamic runtime loader that might cause an application to crash.
 - Corrects a rare illegal instruction system crash.
 - Fixes several race conditions in seg code.
 - Fixes IP multicast packets to work with loopback traffic.
 - Corrects a race condition during disk error recovery failure processing that can result in hung processes
 - Fixes a panic in AIO.
 - Corrects a problem to allow CLSM ioctls to return EINVAL when not in a cluster.
 - Fixes a potential failure during the start up of a multithreaded process.
 - Addresses performance issues seen in select() under certain circumstances.
 - Prevents a potential panic when AdvFS looks up a file name.
 - Fixes a memory fault condition in the user agent driver that occurs when multiple threads issue I/Os on different CPUs.
 - Fixes a number of regular expression problems in multibyte locales and fixes a potential hang problem with complex regular expressions
 - Corrects a problem in which a read or write operation to a changer device creates an unkillable process.
 - Addresses a problem in the IPv6 subsystem where under certain conditions a system may panic.
 - Addresses a scaling issue seen on large multiprocessor systems in dealing with the class scheduling subsystem.
 - Prevents a potential unaligned memory crash when ACLs are on.
 - Corrects a problem of excessive swapping resulting from MFS files system creation.
 - Checks for valid m->m_next in dli_input() to avoid kernel memory fault panic.
 - Corrects a condition in which an error causes a domain panic in bs_frag_dealloc or frag_group_dealloc which results in a panic from a kernel memory fault in access_invalidate and the frag file is not closed properly during the exit from bs_frag_dealloc.
-

Patch 2084.00 Continued

- Fixes a leak of kernel address space.
 - Fixes jitter problems in 24-bit depth on VX1 graphics cards with certain date codes.
 - Fixes a condition in which paging activities on a heavily loaded system can cause the system to crash with a "simple_lock: uninitialized lock" panic.
 - Increases the number of pages that can be pinned at deletion time from 4 to 6. It also improves the informational messages returned by a few domain panic strings.
 - Fixes a problem in the ee driver for DE60x Ethernet adapters
 - Fixes a problem in the PCI bus code that can prevent some functions on a multi-function PCI card from being configured.
 - Fixes a condition in which a system crash can occur if hwmgr deletes a disk while it is in recovery.
 - Fixes a race condition in the kernel AIO code that could panic the system with either a jkernel memory fault or a duplicate malloc free.
 - Fixes a ksh memory fault problem that occurs when logging in.
 - Corrects a problem in which calling aio_write with a negative size may cause a system panic.
 - Corrects problems that can result in changer devices not responding or in a kernel memory fault when the FNDELAY flag is set in changer open and changer ioctl.
 - Adds support for DEGXA-SB and DEGXA-TB Gigabit Ethernet
 - Fixes a memory fault condition in the emx driver that occurs when responding to an inquiry command from a remote port in the fabric.
 - Fixes an I/O hang condition on FibreChannel.
 - Corrects a problem that can cause a core dump when the fwupgrade command is executed with qualifiers.
 - Corrects a Kernel Memory Fault caused by configuring an IPv6 address on a legacy network interface card.
 - Fixes a hole in disk I/O handling that could cause a panic.
 - Eliminates the panic, "lock_terminate: lock owned."
 - Corrects a problem that occurs in some low free memory situations, in which a kernel thread that completes AIO requests may stall on one request, causing other requests (that can complete) to back up behind it.
 - Increases BOOTP error timeout values for Remote Installation Service (RIS) kernel installations.
 - Fixes a potential process hang when exiting that can occur on a system that has dynamically powered off a processor.
 - Fixes an improper handling of domain, area, and fabric RSCN (Registered State Change Notifications) by the emx driver that results in the nondetection of path failures to storage devices in the fabric.
 - Fixes hung I/O in systems with multiple RADs.
 - Corrects a condition that can cause applications using the pshared subsystem to hang. The hang would be seen as the result of a thread failing to wake up after its condition variable timer had expired.
-

Patch 2084.00 Continued

- Updates DDR for the SDLT600 tape drive.
- Fixes panic: simple_unlock: no locks owned by cpu.
- Updates DDR for the DLT VS series of tape devices.
- Addresses issues seen in the taking of non-interactive core dumps using the coredump command.
- Corrects a problem in which the rewind() function would fail to reposition to the beginning of a file.
- Corrects inefficiencies and apparent login hangs involving Enhanced Security in a TruCluster environment.
- Fixes a problem in which a thread can hang during a malloc call.
- Corrects a condition that can cause a system panic caused by stack growth.
- Causes the sysconfigtab settings for the titan_sys_ps_hotswap and titan_sys_fan_hotswap attributes to be maintained after a system reboot.
- Fixes a system hang seen when rename is called with "." as the target and causes other rename argument restrictions to be enforced.
- Corrects a problem in which certain older tape drives (for example, TZK50 and TSV07) would produce errors when they were accessed.
- Fixes possible invalidation of mmap dirty pages before they can be flushed to disk.
- Prevents a potential panic seen with memory_test=partial or memory_test=none.
- Prevents AdvFS metadata inconsistencies and panics on full domain.
- Fixes a problem where an I/O operation may be treated as complete before it has actually finished, thereby possibly causing a panic in XPT or a storage driver.
- Corrects a condition in which active paths may be improperly treated as standby, which could reduce performance or cause a device to switch controllers unnecessarily.
- Fixes a condition in which fork() would return EAGAIN incorrectly on a process using gh_chunks.
- Adds pfil_loopback and pfil_physaddr kernel flags for controlling packetfilter written packets.

Number: Patch 2088.00**Abstract:** Removes the race security vulnerability in find**State:** Supersedes Patches 479.00, 119.00, 451.00, 3.00, 5.00, 447.00, 449.00, 113.00, 115.00, 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, 1040.00, 1576.00, 1577.00, 1583.00, 1940.00, 1941.00, 1943.00, 2059.00

- Corrects a problem in which access to a file may be denied when multiple processes attempt to access the same file at the same time and access to the file should be allowed by an ACL on the file.
 - Corrects a problem in which if the ACL on a file is corrupted, the corrupted ACL is passed into the kernel causing a variety of problems.
 - Corrects a problem in which an AutoFS intercept point for a direct map entry may no longer induce automounts after an error has been detected during a previous automount attempt.
 - Fixes AutoFS problems as follows:
 - Eliminates error messages concerning property lists seen via certain utilities such as vdump.
 - Causes AutoFS automounts 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.
-

Patch 2088.00 Continued

- Prevents a core dump from vdump when a message length is greater than MAX_MSG_SIZE.
- 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.
- Fixes vrestore command problems as follows:
 - Fixes a problem in which the command fails to properly handle extended attributes records in compressed archives. This results in malloc failures, proplist corruption, program abort, program crashes due to segfault or invalid memory access, and the display of the error message "error setting extended attributes"
 - 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 error message "error setting extended attributes."
 - Fixes a problem in which the selective restore of hard-linked 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) option ignores files named "err" but does not ignore files named "errs" and directories named "SCCS" and "RCS".
- Corrects the pax, tar, and cpio commands to properly extract explicitly specified files. When an archive contained a file with extended attributes and a different file (occurring later in the archive) was specified to be extracted, improper buffer pointer management resulted in the following display (example uses tar):

```
tar: /dev/nrmt0h : This doesn't look like a tar archive
tar: /dev/nrmt0h : Skipping to next file...
tar: Memory allocation failed for extended data while reading : Not enough space
```

The directory option was similarly affected. In this case the information for the specified file was not reported.

- Corrects the behavior of several commands when used in conjunction with file systems that are locally served via AutoFS.
 - Provides support for wildcards in Linux /etc/exports entries. Both AutoFS and Automount have been so enhanced.
 - Fixes a problem in which the autofs and autofs mount daemons do not properly parse the asterisk character (*) in map files.
 - Fixes a problem that prevents access to AutoFS file systems if ACLs are enabled.
-

Patch 2088.00 Continued

- Permits 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 the find -ls command so it displays the correct number of blocks.
 - Supports a cluster patch which limits DLM lock contention by AutoFS in a cluster.
 - Fixes a defect that allows the possibility that certain AutoFS auto-mounted file systems may be mounted more than once.
 - Ensures that autofsmount 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.
 - Fixes the following problems with the tar, pax, and cpio commands:
 - Enables tar to check for and report write errors.
 - Gives tar, pax, and cpio the capability to unalter the ctime of input files upon the creation of an archive and enables the display of warning message if pax or cpio is unable to preserve the time of input files.
 - Corrects the behavior of the tar o option.
 - Fixes the cpio -m option in cases where the destination and source files have same mtime.
 - Corrects the pax -l option to create hard links properly.
 - Corrects the cpio -o option so that extended UID file ownership is not corrupted.
 - Fixes the long file names handling in tar.
 - Corrects tar to properly handle unusual directory specifications.
 - Fixes pax to handle ACL on directories properly.
 - Corrects the find -links, -size, -i, -inum behavior with respect to the + operations. Find + operations will match "Greater Than", rather "Greater Than or Equal To."
 - Ensures that tar utility correctly restores directory permissions when extracted using -p option.
 - Fixes an AutoFS panic during an unmount operation when AutoFS tries to remove a directory.
 - Removes the race security vulnerability in find utility.
 - Ensures that tar/pax utilities correctly restore the file mode when extracted using appropriate options.
 - Fixes a problem with the find command that occurs when a large number of directories are searched.
-

TruCluster Server Patches

This chapter provides information about the patches included in Patch Kit 6 for the TruCluster Server software.

This chapter is organized as follows:

- Section 2.1 provides release notes that are specific to the TruCluster Server software patches in this kit.
- Section 2.2 provides brief descriptions of the 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 4 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. References to patch numbers are for TruCluster Server patches unless otherwise indicated.

2.1.1 Required Storage Space

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

- Approximately 250 MB of temporary storage space is required to untar this patch kit (base and TruCluster). We recommend that this kit not be placed in the /, /usr, or /var file systems because doing so may unduly constrain the available storage space for the patching activity.
- Approximately 102 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 104.9 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 2.6 MB of storage space is required in /var/adm/patch/doc for patch abstract and README documentation.
- Approximately 176 KB of storage space is needed in /usr/sbin/dupatch for the patch management utility.

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

2.1.2 Removing Some Patches Can Cause Problems

Removing the following patches from your cluster may cause problems:

- Patches that were installed before a cluster was created.

- Patches that were installed before one or more new members were added to the cluster. This includes running `clu_add_member` during a rolling upgrade after the install stage has been completed.

Because some patches cannot be safely removed in a cluster without causing member-specific problems, we recommend that you do not remove patches under these conditions. Section 2.1.3 describes a problem of this type.

2.1.3 Patch Removal Causes Login Error

If you removed Version 5.1A patches that were installed before a cluster was created or before new members were added, you may see an error similar to the following when you attempt to log into a member:

```
Login Error: Compaq Tru64 UNIX V5.1A (Rev. 1885) (system.xyzcorp.net) console
login:
INIT: Command is respawning too rapidly. Check for possible errors.
id: esmd "/usr/sbin/esmd </dev/null >/dev/null 2>&1
```

If you see this error, remove the following lines from that cluster member's `/etc/inittab` file:

```
esm_init:23:wait:/sbin/init.d/esm init </dev/null >/dev/null 2>&1
esmd:23:respawn:/usr/sbin/esmd </dev/null >/dev/null 2>&1
```

2.1.4 Updates for Rolling Upgrade Procedures

The following sections provide information on rolling upgrade procedures.

2.1.4.1 Problem When Undoing Roll with Worldwide Languages Installed

If, on a system with Worldwide Languages installed, you complete a rolling upgrade of Patch Kit 6 and then run the `clu_upgrade -undo install` command, the tar program may report that it cannot find files it expects to find. This condition is caused by a file left in the `/cluster/admin/tmp` directory from the previous setup stage.

To correct this problem, take the following steps:

1. Undo the setup stage again.
2. Issue the following command:

```
# rm -f /cluster/admin/tmp/*
```

3. Redo the setup stage.

2.1.4.2 Order for Rolling NHD6 and Patch Kit 6

Because Patch Kit 6 contains all of the information contained in the NHD6 kit, you do not need to roll the NHD6 kit in addition to Patch Kit 6. If, however, you plan to roll both kits, you must install the NHD kit first, followed by the patch kit. If you reverse the installation order, you will get a kernel build failure after installing the NHD6 kit.

2.1.4.3 Unrecoverable Failure Procedure

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

2.1.4.4 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.

2.1.4.5 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.

2.1.4.6 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.

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

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

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

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

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

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

```
#!/usr/bin/ksh -p
#
#   Script for restoring installed permissions
#
cd /
for i in /usr/.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)`.

2.1.4.8 Missing Entry Messages Can Be Ignored During Rolling Patch

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

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

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

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

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

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

2.1.4.9 Relocating AutoFS During a Rolling Upgrade on a Cluster

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

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

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

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

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

To minimize your effort in the 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.

2.1.5 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
```

2.1.6 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-0005 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
```

2.1.7 Problems with clu_upgrade switch Stage

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

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

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

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

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

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

2.1.8 Cluster Information for Tru64 UNIX Patch 2084.00

See Section 1.1.20.7 for version switch information related to Tru64 UNIX Patch 2084.00.

2.1.9 Change to gated Restriction — Patch 336.00

The following information explains the relaxed Cluster Alias: `gated` restriction, delivered in TruCluster Patch 336.00.

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.

2.1.10 Version Switch Warning Added — Patch 351.00

TruCluster Server Patch 351.00 provides a warning that informs you that installed patches include a version switch, so those patches cannot be removed using the

normal patch removal procedure. The warning allows you 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.

2.1.11 Information for Patch 371.00

This section provides information for TruCluster Server Patch 371.00.

2.1.11.1 Enablers for EVM

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

2.1.11.2 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*.

2.1.11.3 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.

2.1.12 CAA and Datastore — Patch 338.00

This section provides information about TruCluster Server Patch 338.00.

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 brief descriptions of the patches in Patch Kit 6 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 4)**
Indicates a patch that was new in the specified Version 5.1A patch kit.
- **Existing**
Indicates a patch that was introduced in Patch Kit 1 or Patch Kit 2.
- **Supersedes Patches ...**
Indicates a patch that was combined (merged) with other patches.

Number: Patch 27.00

Abstract: Fix for clusterwide wall messages not being received

State: Existing

- 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

- 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: New

- 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: Patch 136.00

Abstract: Enhancement for clu_autofs shutdown script

State: Existing

- Makes the /sbin/init.d/clu_autofs script more robust.

Number: Patch 188.00

Abstract: Fixes cluster kernel problem that causes a hang

State: Supersedes patches 70.00, 186.00

- 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)

- Fixes a problem in the Memory Channel API that can cause a system to hang.

Number: Patch 212.00

Abstract: Corrects performance issues on starting cluster LSM

State: Supersedes Patch 150.00

- 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.

Number: Patch 254.00

Abstract: Security

State: Supersedes Patch 52.00

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

Number: Patch 256.00

Abstract: Fix for cluster hang

State: Existing (Kit 4)

- 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 265.00

Abstract: Fix for cluster alias manager SUTlet

State: Existing (Kit 4)

- Fixes the problem in which the cluster alias manager SUTlet falsely interprets any cluster alias with virtual={t|f} configured as a virtual alias regardless of its actual setting.

Number: Patch 308.00

Abstract: Corrects various problems with CAA commands

State: Existing (Kit 5)

- Fixes a problem in which some CAA commands, especially caa_profile, in rare scenarios might not function correctly.

Number: Patch 310.00

Abstract: Fixes kernel EVM threads not properly preempting

State: Existing (Kit 5)

- Fixes the potential of multiple assert_wait and timeout panics due to kernel EVM threads not properly preempting.

Number: Patch 314.00

Abstract: Corrects LSM partition types in CNX partition

State: Existing (Kit 5)

- Corrects the LSM partition types in the CNX partition of boot disk for the clu_partmgr utility.

Number: Patch 320.00

Abstract: File names with dollar sign cause upgrade undo probs

State: Existing (Kit 5)

- Fixes rolling upgrade undo problems with file names that contain a dollar sign (\$).
-

Number: Patch 336.00

Abstract: aliasd now interprets NIFF parameters correctly

State: Supersedes Patches 6.00, 7.00, 9.00, 207.00, 208.00, 210.00, 332.00, 333.00, 334.00

- 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.memberX.
 - 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.
 - Fixes a problem in which a core dump occurs in aliasd when all interfaces are removed on a cluster member that is set up with at least one cluster alias that is added with virtual=t and without subnet.
 - Fixes a problem of disabling and reenabling a cluster alias source route on a given interface.
 - Fixes a problem of restoring static routes when interface revives.
 - Corrects a problem in which liasd terminates the incorrect gated process.
-

Number: Patch 338.00

Abstract: Enables CAA to startup and failover system services

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, 242.00, 301.00, 302.00, 304.00

- 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.
 - Fixes a problem in one of the shipped rc scripts whereby Oracle fails during start-up on a clustered system.
 - Fixes the problem that causes the App resource to not go off line when last dependent network resource goes off line.
 - Fixes a problem where CAAD might core dump due to a race condition when multiple events to which it subscribes arrive simultaneously.
 - Fixes the problem that could cause the target member crashes during service start up.
 - Enables CAA to start up and fail over system services before any user services.
-

Number: Patch 351.00

Abstract: Security (SSRT2265)

State: Supersedes Patch 48.00, 138.00, 244.00, 306.00

- Provides a warning to users who have installed a patch kit that includes a patch that requires a version switch. See Section 2.1.10
- Addresses a problem seen during the setup stage of a rolling upgrade, during tag file creation. This patch changes a variable to only look at 500 files at a time, while making tag files instead of the current 700.
- Corrects a potential security vulnerability in the cluster interconnect security configuration that may result in a denial of service.
- Provides clu_upgrade enhancements.
- Corrects a problem with adding a cluster member to a one node cluster during a rolling upgrade. The cluster will no longer not let you advance in the rolling upgrade process.

Patch: Patch 353.00

Abstract: Allow use of 255 in the Lan Interconnect IP address

State: Supersedes Patches 41.00, 80.00, 173.00, 175.00, 246.00

- 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.
 - Corrects a problem to allow the use of "255" in the Lan Interconnect IP address.
-

Number: Patch 355.00

Abstract: Fix rdg: unwiring problem in RDG with local transfers

State: Supersedes Patches 37.00, 82.00, 132.00, 134.00, 182.00, 183.00, 185.00, 249.00, 250.00, 252.00

- 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 uninterruptible 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.
- Fixes multiple problems seen with the TruCluster RDG component. The problems were panics of the following types "rdg: unwiring," "vl_unwire: page is not wired," and "KMF: from _otsmove."

Number: Patch 357.00

Abstract: Security (SSRT2394)

State: Supersedes Patches 50.00, 200.00, 267.00, 269.00, 316.00

- 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.
- Corrects a potential security vulnerability that may result in denial of service. This potential security vulnerability may be in the form of local and remote security domain risks.
- Corrects a problem in which setting the sysconfig inet subsystem values for the tcp_keepcnt attribute to a value < 2 cause the member to panic on boot with the following panic string:

```
NetRAIN configured.  
panic (cpu 0): trap: illegal instruction  
DUMP: Warning: no disk available for dump.
```

- Provides knowledgeable users with the ability to lower the cluster_rebuild_delay.
-

Number: Patch 359.00

Abstract: Correct lock acquires after mpsleep

State: Supersedes Patches 68.00, 257.00, 259.00

- 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."
- Forces a reboot to resolve communications problems in a two node cluster rather than hang.
- Corrects lock acquires after mpsleep.
- Corrects the rebuild delay remainder to be minimally second.

Number: Patch 363.00

Abstract: Fix for ics_mct: RX CONN 3 panic

State: Supersedes Patches 44.00, 46.00, 189.00, 190.00, 191.00, 193.00, 260.00, 261.00, 263.00, 312.00, 360.00, 361.00

- 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 following 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.
 - Corrects an ICS (cluster interconnect) handle memory leak.
 - Addresses a memory leak in the Memory Channel transport layer.
 - Addresses a node hang where a node can become unresponsive as a result of a disconnected Memory Channel cable or a Memory Channel hub being powered down.
 - Addresses an ics_mct: RX CONN 3 panic when a node is joining the cluster.
-

Number: Patch 366.00

Abstract: Fixes a problem in Memory Channel driver

State: Supersedes Patches 11.00, 62.00, 97.00, 145.00, 146.00, 148.00, 203.00, 204.00, 206.00, 273.00, 274.00, 275.00, 277.00, 329.00, 331.00, 364.00

- 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.
 - Facilitates kernel debugging.
 - Corrects a condition that can cause superfluous "rm_event, index too big" messages may appear on a system console.
 - Corrects a problem in a memory channel cluster in which rebooting a node without performing a hardware reset can crash other members with a RM_AUDIT_ACK_BLOCK panic.
 - Fixes issues associated with the initialization of the RM driver.
 - In a memory channel cluster, rebooting a node without performing a hardware reset can crash other members with a RM_AUDIT_ACK_BLOCK panic.
 - Fixes several problems in the Memory Channel driver.
 - Corrects a problem that occurs in a Memory Channel cluster in which rebooting a node without performing a hardware reset can crash other members with an RM_AUDIT_ACK_BLOCK panic.
-

Number: Patch 371.00

Abstract: NFS server with direct I/O assertion panic

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, 272.00, 278.00, 279.00, 280.00, 281.00, 282.00, 283.00, 284.00, 285.00, 286.00, 287.00, 288.00, 289.00, 290.00, 291.00, 292.00, 293.00, 294.00, 295.00, 296.00, 297.00, 298.00, 300.00, 318.00, 321.00, 322.00, 324.00, 325.00, 326.00, 371.00, 339.00, 340.00, 341.00, 342.00, 343.00, 344.00, 345.00, 346.00, 347.00, 349.00, 367.00, 368.00, 369.00

- 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."
 - Addresses 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 enablers for the Enterprise Volume Manager product.
 - Fixes a memory leak in `cfscall_ioctl()`.
 - Provides support for the freezeufs 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 can occur during the mount of a clustered 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`.
-

Patch 371.00 Continued

- 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 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.
 - 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 DMAPI/HSM managed file system:

 (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.
-

Patch 371.00 Continued

- 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.
 - 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 files locked by the plock() function 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.
 - 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 2 GB.
 - 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 (more than one CPU) system.
-

Patch 371.00 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 fail 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.
 - Prevents a “simple_lock: time limit exceeded” panic or an “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.
-

Patch 371.00 Continued

- 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.
 - 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 automounts 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.
 - Fixes a problem in the cluster kernel where a cluster member panics while doing remote I/O over the interconnect.
 - Fixes a performance problem where threads could spend a long time in the check_busy() AdvFS routine.
 - Fixes a panic that may occur during an unmount.
 - Fixes a cross-node cluster deadlock that can occur when AdvFS threads on two cluster nodes simultaneously call code that requires the taking of already held AdvFS locks on the other node.
-

Patch 371.00 Continued

- Corrects a problem in which mounting on a directory in a clone fileset fails with "Device Busy".
 - Enhances cluster file system performance when using file locks to coordinate file access.
 - Prevents a kernel memory fault panic in some cases where AdvFS administration commands are performed on a mounted fileset of an inaccessible AdvFS domain.
 - Fixes a problem in the Device Request Dispatcher.
 - Fixes a race condition in the Device Request Dispatcher.
 - Fixes a problem that occurs when multiple rsh sessions target the cluster alias address, and clua.mod gives out a single port to be used for multiple sessions and cause chaos.
 - Fixes an internal problem in the kernel's AdvFS, UFS, and NFS file systems where extended attributes with names greater than 247 characters could not be set on files. The new limit is 254 plus a null string terminator.
 - Fixes a condition that could cause a panic when a node is halting.
 - Fixes a race condition in CFS readahead logic and a race condition in CFS token logic.
 - Improves the fragment gathering mechanism to boost performance.
 - Fixes a condition that can cause a panic problem when clua.mod is unloaded.
 - Fixes a condition that can cause a boot up panic when ippport_userreserved is 1000 or less.
 - Fixes a problem where access to the quorum disk can be lost if the quorum disk is on a parallel SCSI bus and multiple bus resets are encountered.
 - Fixes a regression associated with non-SCSI storage.
 - Fixes a timing window during asynchronous reads on a CFS client.
 - Fixes a cfsmgr core dump when passing the incorrect number of arguments upon force unmounting a served file system.
 - Fixes a problem in which a CFS client for a file with a hole preceding a frag might drop the frag.
 - Eliminates a performance problem when a node acting as CFS server of an NFS client file system is write-appending to an external NFS server.
 - Fixes a panic that may occur due to a race condition during the mounting of a booting node's boot partition.
 - Fixes a race between nodes performing failover processing which might lead to an incorrect change in the state of file locks.
 - Corrects a problem where, under some circumstances, a system may panic with a kernel memory fault when a device that is being opened by one program is being deleted via the hwmgr utility.
 - Fixes a KMF from mc_bcopy or _OtsMove.
 - Addresses a potential hang in the NFS server when file systems are being relocated in a cluster.
 - Helps to close a race where synchronous writes may obtain disk allocations that were promised to cached client writes.
 - Preserves the error code from an asynchronous write error on a CFS client and returns the error from the close system call.
 - Fixes a potential data inconsistency caused by a problem in the CFS block reservation code which incorrectly calculates the amount of space requested and used by direct I/O writes.
 - Fixes a potential data inconsistency that may occur when a domain is nearly full. The problem causes client write requests shipped synchronously to the server to no longer have subsets of pages written asynchronously due to a race with VM.
-

Patch 371.00 Continued

- Fixes a reconnect problem that occurs when an interface comes down and up again.
- Fixes a panic in clua.mod that is caused by it receiving a delete-cnx request from a member and finding that cnx is in an UNREGISTER state.
- Resolves a kernel memory fault in m_copym.
- Fixes a forced unmount of non-failoverable file system (for example, NFS and AutoFS) panic in the case that the initiator is down.
- Fixes a problem of a panic that occurs in clua.mod when max_aliasid is increased and aliases are added.
- Fixes a rare cluster hang caused by dead locks between the CFS client and server during multiple write operations.
- Fixes an unaligned kernel access in the cluster I/O stack.
- Fixes a problem in which clua.mod does not handle TCP RST messages appropriately.
- Fixes a condition that causes a long delay of an NFS-connection failover when servicing a cluster member that dies.
- Addresses the panic “Assert Failed: (cp->c_flags & CDIRECTIO) == 0” in the cluster file system.
- Causes the correct processing in the close() system call.
- Fixes a condition that causes a boot panic.
- Fixes a condition that causes a CFS client panic during a file system read operation when the server goes down, resulting in the client becoming the server and attempting to release the direction token that had already been released.

Number: Patch 373.00**Abstract:** Fixes problems in the DLM subsystem**State:** Supersedes patches 39.00, 131.00, 178.00, 179.00, 181.00

- 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.
- Corrects a problem in which DLM would not always return the resource block information for the sublock even if the sublock was held.
- Fixes a Distributed Lock Manager panic when calling the dlm_get_lkinfo() routine passing an lkid of a lock block that has already been declared dead by the deadlock detection thread.

Number: Patch 375.00**Abstract:** caa_relocate does not kill the autofsd**State:** New

- Fixes a problem in which caa_relocate AutoFS does not kill the autofsd daemon.
-

Revised Reference Pages

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

A.1 envconfig(8) Update

This section updates the envconfig(8) reference page.

envconfig(8)

NAME

envconfig - Configures the Environmental Monitoring daemon

SYNOPSIS

/usr/sbin/envconfig -c var=value

/usr/sbin/envconfig start | stop

/usr/sbin/envconfig -q

OPTIONS

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

-c var=value

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

ENVMON_CONFIGURED

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

ENVMON_GRACE_PERIOD

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

ENVMON_HIGH_THRESH

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

ENVMON_MONITOR_PERIOD

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

ENVMON_USER_SCRIPT

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

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

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

ENVMON_SHUTDOWN_SCRIPT

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

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

start | stop

Turns the envmond daemon on or off after system startup.

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

DESCRIPTION

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

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

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

EXAMPLES

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

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

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

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

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

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

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

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

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

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

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

FILES

`/etc/rc.config*`

Databases that contains 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

This section updates the `sys_check(8)` reference page.

`syscheck (8)`

NAME

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

SYNOPSIS

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

OPTIONS

`-all`

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

`-debug`

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

`-escalate [ xx ]`

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

with the `-escalate` option to define a crash number.

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

-evm

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

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

-frame

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

-name name

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

-dir name

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

-help or (-h)

Outputs help information.

-nohtml

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

-noquick

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

-perf

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

-v Displays the `sys_check` version number.

-warn

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

-nowarn

Executes only the data gathering pass.

DESCRIPTION

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

The `sys_check` utility also performs an analysis of operating system parameters and attributes such as those that tune the performance of the system.

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

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

Option	Run time	Performance impact	Recommended At
<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 asys_check run. You use the sysconfigdb command or run the shell scripts to implement the stanza files. See the sysconfigdb(8) reference page for more information.

NOTES

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

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

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

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

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

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

The sys_check utility uses other system tools to gather 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 /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 sub-directory, 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)`, `nifftmt(7)`, and `niffconfig(8)` Updates

The installation of Patch 2084.00 results in changes to the `sys_attrs_netrain(5)`, `nifftmt(7)`, `niffconfig(8)`, and `ifconfig(8)` reference pages. Patch 1370.00 updates `ifconfig(8)`; the following sections describe changes to the other reference pages.

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

or equal to its minimum value.

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

SEE ALSO

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

Network Administration: Connections

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:

$$t2 \geq t1 + 2dt$$
$$t1 \geq 0.5$$
$$t2 \geq 1.1$$
$$dt \geq 0.2$$

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

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

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

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

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

    /* tell the traffic monitor to start watching these interfaces */
    while (*av) {
        printf("Adding interface %s to the traffic monitor\n", *av);
        bzero(&mif, sizeof(mif));
        bcopy(*av, &mif.name[0], MIN(strlen(*av)+1, sizeof(mif.name)-1));
        mif.t1 = t1;
        mif.t2 = t2;
        mif.dt = dt;
```

```

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

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

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

```

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.4 wol(8) Update

The installation of Patch 2084.00 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