

HP Tru64 UNIX and TruCluster Server Version 5.1B-6 Patch Summary and Release Notes

Published: October 2010



© Copyright 2010 Hewlett-Packard Development Company, L.P.

Confidential computer software. Valid license from HP 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.

The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

UNIX is a registered trademark of The Open Group. Java is a U.S. trademark of Sun Microsystems, Inc. Oracle is a registered US trademark of Oracle Corporation, Redwood City, California. Linux is a U.S. registered trademark of Linus Torvalds.

Table of Contents

About This Manual.....	11
Audience.....	11
Organization.....	11
Related Documentation.....	11
Patch Process Resources.....	12
HP Encourages Your Comments.....	12
Typographic Conventions.....	12
1 Kit Installation and Removal.....	13
Required Storage Space.....	13
Important Kit Installation and Removal Release Notes.....	13
Installation Release Notes.....	14
Tru64 UNIX NHD-7 Installation Required a Serial Console Connection.....	14
Presence of Some Insight Management Agents Kits May Require Additional Steps (revised).....	14
Some Insight Management Agents Kits May Prevent V5.1B-5 or Higher Installation (revised).....	14
V5.1B-5 or Higher Installation May Overwrite snmpd File (revised).....	15
Stop sendmail Before Installing Kit.....	16
Commands Must Be Run on BIND Systems After Kit Installation.....	16
inetd Daemon Restart Required.....	16
Kit Installation Causes Configuration File Restoration Failure.....	16
Run ipsec Command After Installing Kit.....	17
Procedure to Update lprsetup.dat File.....	17
AdvFS Domain Differences May Affect Version Upgrades.....	17
Possible Errors Seen After Kit Installation.....	18
Message Seen During Reboot Can Be Ignored.....	18
Kit Removal Release Notes.....	18
Some Patch Kits Cannot Be Removed.....	18
Changes to System May Need to Be Reversed.....	19
Script Must Be Run When Returning to Pre-Patched System.....	19
Cluster-Specific Installation and Removal Release Notes.....	20
dupclone Error Message Can Be Ignored.....	20
Installed CSP Could Affect dupatch Cloning Process.....	20
Migrating a Patched Standalone System to a Cluster.....	21
Disable vfast Utility if Running On Cluster Domains.....	21
Creation of Some MFS File Systems Depends on Version Switch.....	21
Workaround Saves Files to Allow Patch Kit Removal.....	21
Enabling the Version Switch After Installation.....	22
Script Required to Reverse Version Switch.....	22
Restriction on Using No-Roll Procedure to Remove Kit.....	22
Do Not Install Prior NHD Kits on a Patched System.....	22
2 Tru64 UNIX Patches.....	23
New Release Notes.....	23
Prior Release Notes.....	23
Dynamic Pathing in Disk Driver.....	23
Performance Improvement for TCP Applications.....	23
Support for 2 TB LUNs.....	23

Enhancements to <code>binary.errlog</code>	23
Updated Printer Support.....	24
ICSNET Pseudo Driver Performance Optimization.....	24
<code>sys_check</code> upgraded to version 145.....	24
Support for the Latest Daylight Saving Time (DST) Changes.....	25
BIND Updated to Version 9.2.8.....	25
Updated Tru64 UNIX Documentation on docs.hp.com.....	25
Standards Conformance.....	26
New Keyword Added to <code>sshd2_config</code> Configuration File for <code>sshd</code> daemon.....	26
New Option to Ignore Processor Set Boundaries.....	26
Support for Evaluating String Comparison Expressions as per POSIX Standards.....	26
<code>iconv</code> Converter Support Surrogate Pairs in Unicode.....	26
New Sysconfig Tunable to Reduce Contention on AdvFS Frag Files.....	27
New <code>rc.config</code> Variables to Hide User Process Arguments and Environmental Variables for <code>ps</code> and <code>w</code> Commands.....	27
Conformance to Open Group Standards.....	27
UNIX 98 Compliance with <code>libc</code>	27
<code>Netstat</code> Read Error on Structures in a Live System.....	27
<code>O_APPEND</code> Flag has no Effect on Behaviour of <code>pwrite()</code>	28
<code>smmsp</code> User and Group Not Required for <code>sendmail</code>	28
Possible Performance slowdown of Oracle 8.1.7 after Tru64 UNIX Rebranding.....	28
Version 5.1B-5 or Higher Kit Requires Uninstallation of Internet Express System Authentication LDAP Module (IAELDAMXXX).....	28
IBM Tivoli Storage Manager (TSM) client problems fixed.....	29
Enhancements Introduced in Prior Kits.....	29
Enhanced Cluster Interconnect Extended to 100 KM.....	29
Cluster Cloning Offers Alternative to No-Roll Patching.....	29
Link Aggregation Extended to Cluster LAN Interconnects.....	29
NetRAIN over LAG Supported.....	30
Support Provided for 2007 Changes to U.S. Daylight Savings Time.....	30
BIND Updated to Version 9.2.5.....	30
Library Calls for Fibre Channel HBA Added.....	31
New Cluster Command Sends ping Packets over TCP.....	32
<code>sendmail</code> Server Updated to Version 8.13.6.....	32
AdvFS Utilities Improved for Working with Metadata Files.....	32
<code>mountd</code> Daemon Gets New Port-Selection Option.....	33
<code>envmond</code> Daemon Modified to Use EVM Events.....	33
<code>aha_chim</code> Driver Problem Corrected.....	33
Command Option Now Provides Additional EMX Driver Information.....	34
New EMX Subsystem Attribute Turns on LLER for Tape I/O.....	34
Kernel Attributes Protect Against ICMP Security Vulnerability.....	34
<code>caa_relocate</code> Command Improved.....	35
<code>collect</code> Utility Improved in Several Ways.....	35
Environment Variable Improves <code>btcreate</code> Kernel Build.....	35
New Variable Aids Performance of AdvFS Administration Commands.....	35
New <code>ftpd</code> Command Option Prevents Login Delays.....	36
New Features Added to <code>kdbx</code> Debugger.....	36
Modified <code>rmvol</code> Utility Allows Multiple Volume Removal.....	36
New <code>disklabel</code> Command Option Expands Partitions.....	37
Commands Modified to Conform to POSIX Standard.....	37
Changes to <code>ex</code> , <code>edit</code> , and <code>vi</code>	37
Changes to <code>awk</code> and <code>nawk</code>	37
Changes to <code>chmod</code>	38
Changes to <code>cp</code>	38
Changes to <code>ex</code>	38

Changes to find.....	38
Changes to rm.....	38
Changes to uucp.....	38
Changes to uuencode.....	39
New Generic Subsystem Attribute Corrects UNIX98 Standards Violations.....	39
waitpid() System Call.....	39
poll() System Call.....	39
New I/O Subsystem Attributes Can Improve Booting Speed.....	40
New Attributes Added to NFS and RPC Subsystems.....	40
New cam Attribute Controls Path Usage.....	41
LSM hot-sparing Improved.....	42
New Option Changes Configuration File Used by aliasd Daemon.....	42
fsdb Utility Now Operates on File System Image.....	42
sendmail Log Problem Corrected.....	42
New Tunable Attribute Corrects NetRAIN Failover Problem.....	42
New Attribute Controls Tape Driver Path Control.....	43
pr Command Behavior Now Works as Described in Reference Page.....	43
Kit's Session Log Made More Useful.....	43
Sys_Check Version 143 Provided.....	43
New Variables Protect Against Attack.....	43
New fuser Option Aids Query Search.....	43
New /etc/printcap Option Provided.....	43
Support for the Name Services Switch Added.....	44
New Hardware Support.....	44
Support for 64 Processor AlphaServer GS1280 Systems.....	44
Support for AlphaServer and AlphaStation DS15 Systems.....	44
HP StorageWorks FCA2384.....	44
Production Version of Motif 2.1 Provided.....	44
Protection Against Buffer Overflow Exploitation Added.....	44
Enhancements to pmgrd Daemon and collect Utility.....	46
Performance Manager Metrics Server Daemon (pmgrd).....	46
New Features Added to collect Utility.....	46
File System Management Applications Enhanced.....	46
New Control Option for /usr/sbin/audit_tool Command.....	47
Change to envmond Improves Performance on Some Systems.....	47
Release Notes Introduced in Prior Kits.....	47
Authentication Choice Affects sftp Transfer Rate.....	47
Tru64 UNIX Rebranding Results in File Changes.....	47
Insight Manager Components DUMP Core.....	47
Autoloader Firmware Upgrade Changes WWND.....	48
Some Smart Array Errors May Not Be Recoverable.....	48
Do Not Use dxarchiver to Verify Bootable Tape.....	48
securenets File Requires localhost Entry.....	48
SIA sialog Use Limitation Required.....	48
Change to executable_data Attribute Requires Running Script.....	48
Potential Security Vulnerability Identified.....	48
Adjusting the tcp_rst_win Variable.....	49
Adjusting the tcp_syn_win Variable.....	49
Modification to Changer Driver May Affect Some Applications.....	50
Data Sorting of Audit Records May Be Required on Single CPU System.....	50
new_wire_method Tunable Attribute Retired.....	50
Stopping Daemons May Speed Administration Performance.....	50
sendmail Application Size/Length Limits Can Cause Problems.....	51
Increasing RDG max_objs Value Recommended.....	51
Reboot May Resolve Problem with Smart Array Controller.....	51

Additional Steps for IPsec Connections.....	51
Potential NFS Duplicate Request Cache Scalability Limitation with High Loads and Uncharacteristic File Access Behavior on Clustered NFS Servers.....	52
Tuning the NFS Server Duplicate Request Cache.....	53
Performance of hwmgr Commands on Large System Configurations.....	54
LSM Spin Lock Issue.....	54
Possible Problem when Processing Many Command Parameters.....	54
Loading Firmware from a BOOTP Server.....	55
Changes to tar, pax, and cpio Behavior.....	55
Changes to vdump and vrestore Disallow Larger Record Sizes.....	55
Problem Seen on Systems with Smart Array Controller.....	56
Broken Links Reported During Baselineing.....	56
Russian Keyboard.....	56
General and Problem Information for AlphaServer ES47, ES80, and GS1280 Systems.....	57
Time Loss on Systems with Firmware Lower Than V6.4-12.....	57
CPU Offline Restrictions.....	57
Problem with Capacity-on-Demand Process.....	57
Repeated Reboots May Cause Panic.....	58
Caution on Updating to Version 5.1B with DEGXA NICs.....	58
Changes to Reference Pages.....	58
Summary of Base Operating System Patches.....	59
New Patches.....	60
Patches Delivered in Previous Kits.....	61
3 TruCluster Server Patches.....	153
New Release Notes.....	153
Prior Release Notes.....	153
New Flag Option to Turn OFF All Existing Flags for Services.....	153
Select Option to Check Tagged Files.....	153
Check for Tagged Files if Messages Are Displayed.....	153
Noncritical Errors.....	154
Unrecoverable Failure Procedure.....	154
Do Not Add or Delete OSF, TCR, IOS, or OSH Subsets During Roll.....	154
Undo Stages in Correct Order.....	154
clu_upgrade undo of Install Stage Can Result in Incorrect File Permissions.....	154
Missing Entry Messages Can Be Ignored During Rolling Patch.....	155
Relocating AutoFS During a Rolling Upgrade on a Cluster.....	155
Messages Displayed During Rolling Upgrade Can Be Ignored.....	156
Error on Cluster Creation.....	156
When Taking a Cluster Member to Single-User Mode, First Halt the Member.....	157
Login Failure Possible with C2 Security Enabled.....	157
Preventing the problem.....	157
Correcting the problem.....	158
File System Unmount Recommended if Message Is Displayed.....	159
Tunable Attribute May Help Performance Problem.....	160
AlphaServer ES47 or AlphaServer GS1280 Hangs When Added to Cluster.....	160
Problems with clu_upgrade Switch Stage.....	161
Data Protector Issues and Restrictions.....	161
Possible Error Backing Up Cluster Mount Points.....	161
Configuring Data Protector for Oracle Integration.....	162
Set ipport_userreserved Attribute on Large Systems.....	162
Summary of TruCluster Server Software Patches.....	162
New Patches.....	162
Patches Delivered in Previous Kits.....	162

4 Worldwide Language Support Patches.....	175
New Release Notes.....	175
Prior Release Notes.....	175
Localization of Message Catalog.....	175
Summary of Worldwide Language Support Patches.....	175
New Patches.....	175
Patches Delivered in Previous Kits.....	175
5 CSPs Included in This Kit.....	179
New Superseded Tru64 UNIX CSPs.....	180
Tru64 UNIX CSPs Superseded in Previous Kits.....	181
New Superseded TruCluster Server CSPs.....	188
TruCluster Server CSPs Superseded in Previous Kit.....	188
A Setting Up an Enhanced Distance Cluster.....	191
Enhanced Distance Cluster Configuration Requirements.....	191
Enhanced Distance Cluster Restrictions.....	192
Enhanced Distance Cluster Recommendations.....	193
B Prior Patch Installation Changes.....	195
Changes Made in Version 5.1B-3.....	195
Changes Made in Version 5.1B-2.....	195
C Component Licensing.....	197
HP Tru64 UNIX Version 5.1B-6 Consolidated Patch Kit.....	197
Apache License.....	202
APACHE HTTP SERVER SUBCOMPONENTS.....	205
mod_mime_magic component.....	205
modules\mappers\mod_imap.c component.....	205
server\util_md5.c.....	205
srclib\apr\include\apr_md5.h component.....	206
srclib\apr\passwd\apr_md5.c Component.....	207
apr_md5_encode() Routine Component.....	207
srclib\apr-util\crypto\apr_md4.c Component.....	207
srclib\apr-util\include\apr_md4.h Component.....	208
srclib\apr-util\test\testdbm.c Component.....	208
srclib\apr-util\test\testmd4.c Component.....	209
srclib\apr-util\xml\expat\conftools\install-sh Component.....	209
PCRE Component.....	210
test\zb.c Component.....	210
expat xml parser Component.....	211
Cyrus IMAP License.....	211
Mozilla License.....	212
OpenLDAP License.....	219
Sendmail License.....	220
Tomcat License.....	221
GNU General Public License.....	224
TCL License.....	229
Perl License.....	230
Zlib License.....	231

List of Figures

A-1	Enhanced Distance Cluster Configuration.....	192
-----	--	-----

List of Tables

2-1	Reference Pages Revised in Previous Version 5.1B Patch Kits.....	58
5-1	Superseded CSPs 02298.01 to 02493.00.....	180
5-2	Superseded CSPs 01936.01 to 02355.00.....	181
5-3	Superseded CSPs 00021.00 to 00283.00.....	182
5-4	Superseded CSPs 00283.01 to 00543.00.....	183
5-5	Superseded CSPs 00545.00 to 00855.01.....	184
5-6	Superseded CSPs 00856.00 to 01107.00.....	185
5-7	Superseded CSPs 00538.01 to 01361.02.....	186
5-8	Superseded CSPs 01362.00 to Patch C 01934.00.....	187
5-9	New Superseded TruCluster Server CSPs.....	188
5-10	TruCluster CSPs Superseded in Previous V5.1B Kits.....	188
5-11	TruCluster CSPs Superseded in Previous V5.1B Kits.....	189
A-1	Hardware Configuration.....	191

About This Manual

This manual contains information specific to Version 5.1B-6 of the Tru64 UNIX® operating system, TruCluster Server software, and Worldwide Language Support products. It briefly describes the patches contained in this kit and provides information you should be aware of when installing or removing this kit.

Audience

This manual is for those who upgrade their operating system to Tru64 UNIX Version 5.1B-6.

Organization

This manual is organized as follows:

Chapter 1 “Kit Installation and Removal”	Provides information you need to be aware of before you install or remove this kit.
Chapter 2 “Tru64 UNIX Patches”	Provides information about the Tru64 UNIX patches included in this kit.
Chapter 3 “TruCluster Server Patches”	Provides information about the TruCluster Server software patches included in this kit.
Chapter 4 “Worldwide Language Support Patches”	Provides information about the patches for the Worldwide Language Support subset included in this kit.
Chapter 5 “CSPs Included in This Kit”	Lists the customer-specific patches (CSPs) for Tru64 UNIX and TruCluster Server that have been superseded by patches included in this kit.
Appendix A “Setting Up an Enhanced Distance Cluster”	Provides information about setting up and configuring an Enhanced Distance Cluster.
Appendix B “Prior Patch Installation Changes”	Provides a brief overview of important changes made to the patch kit installation process introduced in previous patch kits.
Appendix C “Component Licensing”	Provides the licenses for software components included in the Tru64 UNIX Version 5.1B-5 kit.

Related Documentation

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

- *Technical Updates for the Version 5.1B and Higher Operating System and Patches*
This online supplement reports any information about restrictions and problems that may have been discovered since the release of the Version 5.1B operating system and its patch kits.
<http://h30097.www3.hp.com/docs/updates/V51B/html/index.html>
- *Patch Kit Installation Instructions*
<http://h30097.www3.hp.com/docs/patch/install/HTML/TITLE.HTM>
- *Patching Best Practice*
http://h30097.www3.hp.com/docs/best_practices/BP_PATCH/TITLE.HTM
- 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.
<http://h30097.www3.hp.com/docs/patch/dupatch.8.html>
- The following Tru64 UNIX and TruCluster Server software installation and administration guides:

- [Tru64 UNIX Installation Guide](#)
- [Tru64 UNIX System Administration](#)
- [TruCluster Server Cluster Installation](#)
- [TruCluster Server Cluster Administration](#)

These guides, and most of the documentation related to the Tru64 UNIX operating system, are available online at the following Web Page:

<http://h30097.www3.hp.com/docs/>

- Release-specific installation documentation

Patch Process Resources

We provide websites to help you with the patching process:

- To obtain the latest patch kit for your operating system and cluster:
<http://www2.itrc.hp.com/service/patch/mainPage.do>
- To obtain early release patches (ERPs):
<http://h30097.www3.hp.com/unix/EarlyReleasePatch-download.html>
- 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/>
- To visit our Business Support Center:
<http://h20000.www2.hp.com/bizsupport/TechSupport/Home.jsp>
- To visit the Tru64 UNIX homepage:
<http://h30097.www3.hp.com/>

HP Encourages Your Comments

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

You can send your comments using the following website: <http://h30097.www3.hp.com/comments.html>

Typographic Conventions

This document uses the following typographical conventions:

#	A number sign represents the superuser prompt.
<i>audit(5)</i>	A reference page. The reference page name is <i>audit</i> , and it is located in Section 5.
Command	A command name or qualified command phrase.
Computer output	Text displayed by the computer.
User input	Commands and other text that you type.
<i>Variable</i>	The name of a placeholder in a command, function, or other syntax display that you replace with an actual value.
device names	Operating system versions before Version 5.0 use different names than those of Version 5.0 and higher. In general, this manual uses the Version 5.0 names. For example, where a partition name is represented by <code>/dev/disk/dsk3g</code> , the Version 4.0n name might be <code>/dev/rz3g</code> .

1 Kit Installation and Removal

This chapter provides information you need to be aware of before you install or remove this kit. It is organized as follows:

- The “[Required Storage Space](#)” section lists the approximate storage space requirements for this patch kit when installing the operating system patches alone and in combination with the TruCluster Server and Worldwide Language Support (WLS) patches.
- The “[Important Kit Installation and Removal Release Notes](#)” section describes issues you need to be aware of when installing or removing this kit.

If you have not installed Version 5.1B-3 or earlier kits, see [Appendix B “Prior Patch Installation Changes”](#) for changes to the `dupatch` utility that you should be aware of before installing this kit.

Required Storage Space

Approximately 250 MB of temporary storage space is required to untar this patch kit. It is recommended that you do not place this kit in the `/`, `/usr`, or `/var` file systems; doing so may unduly constrain the available storage space for the patching activity.

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

- The following approximate amount of storage space in `/var/adm/patch/backup` may be required for archived original files if you choose to install and revert all patches.
 - 108.4 MB for Tru64 UNIX alone
 - 125.8 MB for Tru64 UNIX and TruCluster Server
 - 111.4 MB for Tru64 UNIX and WLS subset
 - 132.1 MB for Tru64 UNIX, TruCluster Server, and WLS subset
- The following approximate amount of storage space in `/var/adm/patch` may be required for original files if you choose to install and revert all patches:
 - 105.2 MB for Tru64 UNIX alone
 - 122.3 MB for Tru64 UNIX and TruCluster Server
 - 108.1 MB for Tru64 UNIX and WLS subset
 - 139.5 MB for Tru64 UNIX, TruCluster Server, and WLS subset
- The following approximate amount of storage space in `/var/adm/patch/doc` may be required for the patch abstract and README documentation:
 - 3.1 MB for Tru64 UNIX alone
 - 3.3 MB for Tru64 UNIX and TruCluster Server
 - 3.1 MB for Tru64 UNIX and WLS subset
 - 3.4 MB for Tru64 UNIX, TruCluster Server, and WLS subset
- Approximately 200 KB of storage space is needed in `/usr/sbin/dupatch` for the patch management utility.

For more information, see the [Patch Kit Installation Instructions](#).

Important Kit Installation and Removal Release Notes

The release notes in this section provide information you need to be aware of when installing or removing this kit. Notes indicated as “(new)” do not appear in the release notes that shipped with prior 5.1B versions. Notes indicated as “(revised)” were included in the previous kit but revised for this release.

Also be aware of any Special Instructions that may appear on screen when running the `dupatch` program.

The following topics are addressed:

- “Installation Release Notes” (page 14) provides information about installing this kit.
- “Kit Removal Release Notes” (page 18) provides information about removing this kit.
- “Cluster-Specific Installation and Removal Release Notes” (page 20) provides information specifically for the installation and removal of this kit on a system running the TruCluster Server software.

Installation Release Notes

The following notes provide important information you need to be aware of before installing the Version 5.1B-6 kit.

Tru64 UNIX NHD-7 Installation Required a Serial Console Connection

HP recommends the following procedure for installing the DS-A5134-AA in an Alpha System GS1280:

1. Install Tru64 UNIX NHD-7 – must be installed with a serial console connection.
2. Install and configure HBA (Host Bus Adapter) DS-A5134-AA and/or reconfigure boot disks partitions – use either a graphics console or a serial console connection.

Presence of Some Insight Management Agents Kits May Require Additional Steps (revised)

The following installation-related release notes pertain to the Insight Management Agents.

It is strongly recommended that any existing version of Tru64 UNIX Insight Management Agents kit (CQPIMxxx kit, where xxx=310, 320, 370, and so on) be uninstalled prior to Version 5.1B-5 or higher update and CPQIM370 kit with latest available CPQIM370 patches be re-installed after Version 5.1B-5 or higher update. The CPQIM kit is available at: <http://h30097.www3.hp.com/cma/>

The CPQIM patches are available at:

<http://h30097.www3.hp.com/cma/patches.html>



NOTE: See the section “Insight Manager Components DUMP Core” (page 47) for information about a potential problem with Insight Management Agents that can occur after Version 5.1B-5 or higher is installed.

Some Insight Management Agents Kits May Prevent V5.1B-5 or Higher Installation (revised)

Under certain conditions, you will be prevented from installing Version 5.1B-5 or higher if you are running HP Tru64 UNIX Insight Management Agents Version 3.1 or higher or had a version of the kit previously installed. Those conditions are as follows:

- Your system contains a pre-Version 5.1B-5 or higher kit and the Insight Management Agents kit.

In this case, upgrading to this kit gives the following error message:

```
Patch 29020.00 - SP07 OSFCLINET540 (SSRT5971 SSRT3653 SSRT2384 ...)
./sbin/init.d/snmpd: its origin can not be identified.
This patch will not be installed.
```

- Your system contains Patch Kit 2, Patch Kit 3, or Patch Kit 4 and the Insight Management Agents kit was once installed but has since been removed.

In this case, upgrading to Version 5.1B-5 or higher gives the following error message:

```
Patch 29020.00 - SP07 OSFCLINET540 (SSRT5971 SSRT3653 SSRT2384 ...)
./etc/pmgrd_iorate.config: does not exist on your system,
however, it is in the inventory of installed subsets.
This patch will not be installed.
```

To work around this problem you will need to run the `dupatch` baseline process before installing Version 5.1B-5 or higher. The following steps will guide you through the process:

1. Create a backup copy of the `/sbin/init.d/snmpd` script. For example:

```
# cp /sbin/init.d/snmpd /tmp
```

An alternative to backing up this file, in which you manually modify it, is provided following step 7.

2. Run the Version 5.1B-5 or higher `dupatch` utility and select Option 5, Patch Baseline Analysis/Adjustment. For detailed instructions, see the [Patch Kit Installation Instructions](#).

3. After Phase 5 of the baseline procedure, answer `y` to the following question:

```
Do you want to enable the installation of any of
these patches? [y/n]: y
```

Phase 5 reports patches that do not pass installation applicability tests due to the current state of your system. The installation of Patch 29020.00 was prevented because of changed system files. The `dupatch` utility reports the known information about the files contained in each patch and asks if you want to enable the installation. Answering yes, enables `dupatch` to install patches that were prevented from being installed due to unknown files.

4. Install Version 5.1B-5 or higher.
5. After the system is running with 5.1B-5 or higher installed, stop the `snmpd` and `insightd` daemons as follows:

```
# /sbin/init.d/snmpd stop
# /sbin/init.d/insightd stop
```

6. Replace the `/sbin/init.d/snmpd` script with the one you copied in step 1; for example:

```
# cp /tmp/snmpd /sbin/init.d/snmpd
```

7. Start the `snmpd` and `insightd` daemons as follows:

```
# /sbin/init.d/snmpd start
# /sbin/init.d/insightd start
```

If you did not back up the `/sbin/init.d/snmpd` file in step 1, you can modify it after you install Version 5.1B-5 or higher (step 4) and stop the `snmpd` and `insightd` daemons (step 5) as follows (the XXX represents the revision, such as CPQIM360):

1. Edit the line that reads `CPQMIBS=/usr/sbin/cpq_mibs` as follows:

```
CPQMIBS=/var/opt/CPQIMXXX/bin/cpq_mibs
```

2. Edit the line that reads `PMGRD=/usr/sbin/pmgrd` as follows:

```
PMGRD=/var/opt/CPQIMXXX/bin/pmgrd
```

3. Edit the line that reads `$PMGRD > /dev/console 2>&1 &` as follows:

```
$PMGRD ` $RCMGR get PMGRD_FLAGS ` > /dev/console 2>&1 &
```

V5.1B-5 or Higher Installation May Overwrite snmpd File (revised)

When you install a newer version of the Insight Management kit, the paths to the `cpq_mibs` and `pmgrd` subagents are changed in the `snmpd` script. By installing Version 5.1B-5 or higher, the `snmpd` script is replaced by the original version provided in the base version of the Insight Management kit.

Because the use of that `snmpd` script will cause problems when using Insight Manager, you must restore the script to the latest version. To do this, create a backup file of the `snmpd` script and restore the backup version after installing V5.1B-5 or higher. (See step 1 of the workaround described in “Some Insight Management Agents Kits May Prevent V5.1B-5 or Higher Installation (revised).”)

If you did not back up the `snmpd` file before installing V5.1B-5 or higher, you can modify the file after the installation, as described in “Some Insight Management Agents Kits May Prevent V5.1B-5 or Higher Installation (revised).”

Stop sendmail Before Installing Kit

It is important that you stop the `sendmail` mailer daemon before installing this kit. Failing to do so can lead to the loss of queued mails. Lost mails cannot be recovered.

To stop the daemon, enter the following command:

```
# /sbin/init.d/sendmail stop
```

Commands Must Be Run on BIND Systems After Kit Installation

After installing this kit on a system configured to be BIND server, run the following command:

```
# rcmgr set BIND_SERVERARGS "-c /etc/namedb/named.conf"
```

On a cluster configured to be a BIND server, run the following command:

```
# rcmgr -c set BIND_SERVERARGS "-c /etc/namedb/named.conf"
```



NOTE: Note that in BIND 9, the `named` daemon uses the `c` option to pass a configuration file parameter, instead of the `C` option that was used in previous versions of BIND.

Stop the `named` daemon and restart it in order to have the new `named` daemon take effect:

- For standalone systems:

```
# /sbin/init.d/named stop
# /sbin/init.d/named start
```

- For clusters:

```
# /sbin/init.d/named cluster_stop
# /sbin/init.d/named start
```

To verify that your configuration files are compatible with Bind 9, run the following commands:

```
# named-checkconf /etc/namedb/named.conf
# named-checkzone example.com /etc/namedb/hosts.db
```



NOTE: With BIND 9, CNAME entries no longer accept quotes. For example, `"hosts-1" IN CNAME A` needs to be changed to `hosts-1 IN CNAME A`

See “BIND Updated to Version 9.2.8” (page 25) for information about BIND 9.

inetd Daemon Restart Required

Because of changes made to the Internet services daemon introduced in this release, you need to stop and then restart `inetd` after installing or removing this kit. You can do this from the command line or by using the `sysman` application. From the command line, enter the following commands:

```
# /sbin/init.d/inetd stop
# /sbin/init.d/inetd start
```

Failure to do this will result in an older version of `inetd` running on your system.

Kit Installation Causes Configuration File Restoration Failure

After installing this kit, attempts to restore the configuration file (`config.cdf`) saved prior to the installation of this patch will fail due to a checksum error. You can, however, force the operation by using the following `sysman` command:

```
# sysman -clone -apply -force config.cdf
```

For more information, see the note titled *Correction to Configuration Cloning Restrictions* in the “Corrections to Manuals” section of the online Technical Updates document for Version 5.1B. The following link will take you to the Technical Updates document:

<http://h30097.www3.hp.com/docs/updates/V51B/html/index.html>

Run ipsec Command After Installing Kit

If you are running IP Security (`ipsec`) on your system, run the following command after installing this kit to determine if any unsafe connections exist:

```
# /usr/sbin/sysman ipsec
```

A warning message will alert you to any potential problems.

Procedure to Update lprsetup.dat File

If you use the `/usr/sbin/printconfig` application to configure printer queues, run the following command as root to update the `/etc/lprsetup.dat` file:

```
# /usr/sbin/lprsetup -c update
```

AdvFS Domain Differences May Affect Version Upgrades

A difference in the structure of Version 5.1A and early 5.1B AdvFS domains versus later V5.1B domains can cause a problem when upgrading to Version 5.1B-4 or higher.

This potential problem involves a metadata file called the RBMT that exists on each volume of a version 4 domain.

Although an RBMT is generally only one page long, it can be longer on large volumes or domains that have many files. If an RBMT file was larger than one page long under 5.1A or an early 5.1B version and then grows again after a system upgrade to Version 5.1B-4 or higher, the RBMT file can cause a problem in which any command that tries to activate that domain will fail. This includes mounting filesets from the affected domain.

Following a system upgrade to Version 5.1B-4 or higher, the problem can occur after all the filesets in a domain are unmounted. (The problem will not occur as long as the filesets remain mounted.)

The solution is to use the `fixfdmn` utility to correct the problem. For example:

```
# /sbin/advfs/fixfdmn domain_name
fixfdmn: Checking the RBMT.
fixfdmn: Clearing the log on volume /dev/disk/dsk10c.
fixfdmn: Checking the BMT mcell data.
fixfdmn: Checking the deferred delete list.
fixfdmn: Checking the root tag file.
fixfdmn: Checking the tag file(s).
fixfdmn: Checking the mcell nodes.
fixfdmn: Checking the BMT chains.
fixfdmn: Checking the frag file group headers.
fixfdmn: Checking for frag overlaps.
fixfdmn: Checking for BMT mcell orphans.
fixfdmn: Checking for file overlaps.
fixfdmn: Checking the directories.
fixfdmn: Checking the frag file(s).
fixfdmn: Checking the quota files.
fixfdmn: Checking the SBM.
fixfdmn: Completed.
```

You can use this command proactively before the RBMT grows to prevent the problem from occurring or you can use it after the problem occurs.

In summary, the following domains are not in danger:

- Version 3 domains
- Domains created under Version 5.1B-4 or higher
- Domains with RBMT files that are not longer than one page

The `showfile` and `showdmn` commands can provide information about your domains.

Use the `showdmn` command to find out what volumes a domain has. For example:

```
# /sbin/showfdmn domain_name
      Id          Date Created  LogPgs  Version  Domain Name
447350cd.000eba90  Tue May 23 11:13:33 2006      512      4  domain_name
Vol  512-Blks      Free  % Used  Cmode  Rblks  Wblks  Vol Name
  1L  71132000      71121632      0%    on    256    256  /dev/disk/dsk4c
```

Use the `showfile` command to determine if an RBMT file has more than one page. To do this, select any mounted fileset from the domain in question, find the mount point for the fileset, and enter the following command. (Note that `.tags/M-6` represents volume 1. Subsequent volumes are incremented by a factor of six, so that volume 2 uses `.tags/M-12`, volume 3 uses `.tags/M-18`, and so on.) For example:

```
# /usr/sbin/showfile mountpoint/.tags/M-6
      Id  Vol  PgSz  Pages  XtntType  Segs  SegSz  I/O  Perf  File
fffffffa.0000  1  16    1    simple   **    **   ftx  100%  M-6
```

See the `fixfdmn(8)`, `showfile(8)`, and `showfile(8)` reference pages for information about using these commands.

Possible Errors Seen After Kit Installation

The following problems have been known to occur after Version 5.1B-4 or higher has been installed:

- The Common Data Security Architecture (CDSA), IP Security Protocol (IPsec), or Single Sign-On (SSO) do not work.
- The following error message is displayed during boot time:

```
CSSM_ModuleLoad: CSSM error 4107
```

If you experience these problems, make sure that the following command line has been executed:

```
# /usr/sbin/cdsa/mod_install -f -i -s \
/usr/lib/cdsa/libt64csp.so -d /usr/lib/cdsa/
```

Message Seen During Reboot Can Be Ignored

The following error message will be displayed after you reboot your system the first time after installing Version 5.1B-4 or higher:

```
AllowCshrcSourcingWithSubsystems is not valid
ForcePTYAllocation is not valid
IdentityFile is not valid
AuthorizationFile is not valid
```

These messages are caused by a new version of SSH included in Version 5.1B-4 or higher. They do not pose a problem and can be ignored.

Kit Removal Release Notes

The following sections describe actions you have to take if you decided to uninstall Version 5.1B-4 or higher. Read each section before running the patch deletion procedure.

Some Patch Kits Cannot Be Removed

You cannot remove a patch kit on systems that have New Hardware Delivery 7 (NHD-7) kit installed when either of the following conditions exist:

- The patch kit you want to remove was installed before the NHD (New Hardware Delivery) kit.

For example, if you installed Patch Kit 2 and then installed NHD-7, you cannot remove that patch kit. However, if you later installed Patch Kit 4, you can remove that patch kit.

- The patch kit was installed with NHD-7.

Beginning with the release of Patch Kit 3, patch kits were incorporated into the NHD-7 kits. As a result, when you installed NHD-7, you automatically installed the current patch kit.

These patch kits cannot be removed. However, you can remove any subsequent patch kits. For example, if you installed NHD-7 with Patch Kit 4 and later installed Patch Kit 5, you cannot remove Patch Kit 4, but can remove Patch Kit 5.

If you must remove the patch kit, the only solution is to rebuild your system environment by reinstalling the Version 5.1B operating system and then restoring your system to its state before you installed NHD-7 with the unwanted patch kit.

Changes to System May Need to Be Reversed

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

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

To uninstall this kit, do the following:

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

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

Script Must Be Run When Returning to Pre-Patched System

If removing this patch kit restores your system to a pre-patched state, you must run the script `/etc/dn_fix_dat.sh` before rebooting your system during the patch-deletion process.

This situation would occur if Version 5.1B-2 or higher is the only Tru64 UNIX patch kit installed on your 5.1B system.



NOTE: Because the no-roll procedure automatically reboots the system after deleting the patches, you cannot use this method to delete this patch kit if doing so returns your system to a pre-patched state.

Failing to run this script will result in your system being unable to boot normally. If this occurs, do the following:

1. Boot your system in single-user mode:

```
>>> boot -fl s
```

2. Run the script:

```
# /etc/dn_fix_dat.sh
```

3. Reboot normally.

If you also need to reverse the version switch as described in “Script Required to Reverse Version Switch”, run the `/etc/dn_fix_dat.sh` script after step 5 in that process.

Because the no-roll procedure automatically boots your system, you cannot use that patch kit removal method if doing so would restore your system to a pre-patched state



Note:

If during the `dupatch` installation and deletion processes you see a Special Instruction about running this script, you should ignore that instruction unless your system meets the requirement described here.

Cluster-Specific Installation and Removal Release Notes

This section provides information you need to be aware of if you are installing or removing patch kits from a TruCluster Server environment.

dupclone Error Message Can Be Ignored

Installing this kit using the `dupclone` process on systems that do not have all of the operating system and TruCluster Server base subsets installed may result in a messages similar to the following to be displayed:

Problem installing:

- Tru64_UNIX_V5.1B:
Patch 29034.00

requires the existence of the following un-installed/un-selected subset(s):

- Tru64_UNIX_V5.1B:
Patch 29023.00
- Tru64_UNIX_V5.1B:
Patch 29050.00

You can ignore this message. In all cases, the subsets will be installed correctly.

See “Cluster Cloning Offers Alternative to No-Roll Patching” (page 29) for an introduction to `dupclone`.

Installed CSP Could Affect dupatch Cloning Process

If you have installed customer-specific patches (CSPs) on your system, you may see a message similar to the following when installing this kit using the `dupatch` cloning process, at which time the cloning process will be terminated:

```
Inspecting 69 patches for possible system conflicts ...
./usr/bin/ls:
    is installed by Customer Specific Patch (CSP):

- Tru64_UNIX_V5.1B / Installation Patches:
  Patch 01682.00 - Fix for dupatch command and can not be replaced by this patch
                  . To install this patch, ideally, you must first remove the
                  CSP using dupatch. Before performing this action, you should
                  contact your HP Service Representative to determine if this
                  patch kit contains the CSP. If it does not, you may need to
                  obtain a new CSP from HP in order to install the patch kit
                  and retain the CSP fix. Or, you may use dupatch baselining to
                  enable the patch installation.
```

The recommended action is to perform `dupatch` baselining on your existing system to enable the patch installation process and retain the CSP on your system. Removing the CSP (as mentioned in message) could eliminate the fixes made by that CSP.

After running the baselining process on your existing system, you will need to begin the cloning process from the beginning by reduplicating your system on an alternate set of disks and rerunning the `dupatch` cloning process. See the [Patch Kit Installation Instructions](#) for information about performing baselining and on the patch cloning process.

Migrating a Patched Standalone System to a Cluster

Installing only the base patches on a non-cluster system omits various patches (including some security patches) because of dependencies on TruCluster Server patches. Such patches are not needed in standalone systems. However, if the standalone system is then clustered using the `clu_create` command and you attempt to apply the cluster patches, many patches will fail with errors because various prerequisite patches failed.

These errors do not necessarily indicate that the patch process has failed, but they are numerous, can be confusing and might obscure genuine errors.

The preferred procedure for adding a standalone system into a cluster is as follows:

1. Reinstall the operating system on the standalone system.
2. Run the `clu_create` command and bring the standalone system as a cluster node.
3. Apply all base and cluster patches.

Disable vfast Utility if Running On Cluster Domains

If the `vfast` utility is running on the TruCluster domains `cluster_root` and `cluster_var`, deactivate it on the domains before installing or removing this kit. To deactivate `vfast` on the two domains, use the following command:

```
# vfast deactivate cluster_root
# vfast deactivate cluster_var
```

See the `vfast(8)` reference page for more information.

Creation of Some MFS File Systems Depends on Version Switch

During the installation of this kit, MFS file systems that are 4 GB and larger (or 2 GB and larger if a 1024-byte sector size is used) cannot be created until after the version switch is thrown. (See the [Patch Kit Installation Instructions](#) for information about the version switch.)

Workaround Saves Files to Allow Patch Kit Removal

If you upgrade the operating system and install a patch kit within the same roll, the contents of the patch backups are inadvertently removed. The result is that the patches most recently installed cannot be removed because the backups are missing.

The following procedure saves then restores backups so they will be available if you later decided to remove the patch kit:

1. Create backup files of the `/backup` and `/doc` directories after the postinstall step (`clu_upgrade postinstall`) as follows:

```
# cd /var/adm/patch/backup
# tar cvf /var/adm/patch/BACKUP.tar *
# cd /var/adm/patch/doc
# tar cvf /var/adm/patch/DOC.tar *
```

2. After the switch step (`clu_upgrade switch`) untar the files you created in step 1:

```
# cd /var/adm/patch/backup
# tar xvf /var/adm/patch/BACKUP.tar
# cd /var/adm/patch/doc
# tar xvf /var/adm/patch/DOC.tar
```

This will restore the files under the following directories:

- `/var/adm/patch/backup`
- `/var/adm/patch/doc`

Enabling the Version Switch After Installation

Some patches require you to run the `versw -switch` command to enable the new functions delivered in those patches. (See the [Patch Kit Installation Instructions](#) for information about version switches.) Enter the command as follows after `dupatch` has completed the installation process:

```
# versw -switch
```

The new functionality will not be available until after you reboot your system. You do not have to run the `versw -switch` command, but if you do not, your system will not be able to access the functionality provided in the version-switch patches.

Script Required to Reverse Version Switch

If you enabled version switches as described in the section titled “Enabling the Version Switch After Installation”, you must run the `/usr/sbin/versw_enable_delete` script before attempting to remove Version 5.1B-4 or higher. The steps for running this script require a complete cluster or single system shutdown, so choose a time when a shutdown will have the least impact on your operations. The following steps describe the procedure:

1. Make sure that all phases of the installation process have been completed.
2. Run the `/usr/sbin/versw_enable_delete` script:

```
# /usr/sbin/versw_enable_delete
```
3. Shut down the entire cluster or the single system.
4. Reboot the entire cluster or the single system.
5. Run `dupatch` on your single system or on a cluster using the rolling upgrade procedure to delete Version 5.1B-4 or higher (as described in the [Patch Kit Installation Instructions](#)), up to the point where the kernel is rebuilt and the system must be booted.
6. Reboot the single system or each member of the cluster.



NOTE: This step requires that you reboot each cluster member to remove Version 5.1B-4 or higher. Because the no-roll procedure automatically reboots the system after deleting the patches, you would not be able to perform this step as required.

Restriction on Using No-Roll Procedure to Remove Kit

The section titled “Script Must Be Run When Returning to Pre-Patched System” describes actions you need to take before rebooting your system if removing this kit would restore your system to a pro-patched state. Because the no-roll procedure automatically boots your system, you cannot use that patch kit removal method if doing so would restore your system to a pre-patched state.

Do Not Install Prior NHD Kits on a Patched System

Do not install the NHD-5 or NHD-6 kits on your TruCluster Server system if you have installed this patch kit or earlier patch kits. Doing so may cause an incorrect system configuration. The installation code for these new hardware delivery kits does not correctly preserve some cluster subset files.

2 Tru64 UNIX Patches

This chapter provides information about the defect fixes included in Version 5.1B-6 for the base operating system.

This chapter is organized as follows:

- The “New Release Notes” section lists release notes that are specific to the Tru64 UNIX patches in this kit, as well as release notes that are of general interest.
- The “Prior Release Notes” section lists release notes listed from the initial Version 5.1B release through Version 5.1B-5.
- The “Changes to Reference Pages” section lists reference pages that have been updated or added since the initial Version 5.1B release.
- The “Summary of Base Operating System Patches” section provides brief descriptions of the changes delivered in this patch kit and in prior Version 5.1B patch kits.

New Release Notes

This release contains only defect fixes. There are no new features or enhancements added in this release of Tru64 UNIX software.

Prior Release Notes

Because patch kits are cumulative, this kit will install all of the fixes, features, and changes that have been added since you last installed a Version 5.1B patch kit. The following sections describe the changes contained in this kit that were first introduced in prior kits.

Dynamic Pathing in Disk Driver

This kit enhances the CAM disk driver to dynamically recognize and use newly added paths while the device is in use.

The new behavior will now dynamically recognize and use new paths while the device is in use, thereby making all the paths available for I/O.

Before Version 5.1B-5, the disk driver used all paths that were available at open/mount time. Paths added after would not be added to the active set until a close/open or dismount/mount occurred.

Performance Improvement for TCP Applications

The TCP selective acknowledgment feature is enabled by default. When the selective acknowledgments feature is enabled, the data receiver can inform the sender about all the segments that were received successfully, so that the sender needs to retransmit only those segments that were lost. This improves the performance of TCP applications on a network which is experiencing packet loss.

Support for 2 TB LUNs

This kit enhances the Tru64 UNIX CAM (Call Applications Manager) subsystem to support a LUN size of upto 2 TB from the previous maximum limit of 1 TB. This allows Tru64 UNIX to take advantage of the increased LUN size supported by the storage arrays.

Enhancements to `binary.errlog`

This kit enhances the error log feature to improve user experience. In some error situations, the entries appear to be missing time sequences within the `binary.errlog` file, which makes it difficult to determine if the system did not have any events to log or there was some issue (such as low disk free space) that disabled event reporting for a period of time.

This kit introduces the following additional enhancements to `binary.errlog` file:

1. Introduce a set of markers in the `binary.errlog` file frame work to track the events occurring in the CAM layer and log the most recent log that was attempted on the system.
2. Increase the internal buffer `.blbuf` through the `sysconfigtab` variable, if necessary.
3. Notify the user upon disk full state to clear the disk space and restart the `binlogd` daemon.
4. Provide crash extensions to dump `.blbuf` and the newly introduced track framework data structures.

Updated Printer Support

This kit introduces support for the following 18 printers:

HP LaserJet 1300	HP LaserJet 9050 MFP
HP LaserJet 1320	HP LaserJet 9055 MFP
HP LaserJet 2410	HP ColorLaserJet 3000
HP LaserJet 2420	HP ColorLaserJet 3700
HP LaserJet 2430	HP ColorLaserJet 3800
HP LaserJet 4345 MFP	HP ColorLaserJet 4730 MFP
HP LaserJet 5200	HP ColorLaserJet 5500
HP LaserJet 9000 MFP	HP ColorLaserJet 5550
HP LaserJet 9040 MFP	HP ColorLaserJet 9500 MFP

ICSNET Pseudo Driver Performance Optimization

This kit provides an enhancement to the ICSNET pseudo driver that improves the performance of the `ics0` interface in a LAN cluster.

This optimization is helpful for cluster applications that use the `ics0` interface to interact with other nodes in the cluster. It does this by avoiding the latency associated with time critical Interconnect Communication Subsystem (ICS) remote procedure calls, and by using available bandwidth of the LAN interconnect directly, which provides increased throughput.

This new functionality is enabled or disabled by the new `sysconfigtab` attribute, described as follows:

```
icsnet:  
icsnet_optimization_enable = 1
```

The following restrictions and limitations apply to the use of this feature:

1. This enhancement requires the version switch. Run the `/var/adm/patch/noroll/noroll_versw` command, after the no-roll installation.
2. There will not be any change in the statistics of the `ics0` interface. All statistics must be checked at the physical cluster interconnect level only.
3. `Tcpdump` behavior will change. The sender side behavior remains the same. However, the receiving part of `tcpdump` must be checked at the physical interface (cluster interconnect) level.
4. The `ics0` interface MTU (Magnetic Tape Unit) will be dependent on the MTU of the physical cluster interconnect. Any change in the MTU of the physical cluster interconnect without a reboot, requires a change in the `ics0` interface's MTU as well.
5. The ICSNET optimization feature is not supported in configurations where cluster interconnects are configured as VLAN (virtual LAN).

`sys_Check` upgraded to version 145

This kit includes `sys_Check` Version 145. However, HP recommends that you visit the `sys_check` website to download and install the latest version of `sys_check`:

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

Support for the Latest Daylight Saving Time (DST) Changes

This kit updates the `/etc/zoneinfo` time zone data files to incorporate all changes as of (date of change) in time zones around the world, most notably the following:

Australia

The Australian provinces of Victoria, New South Wales, South Australia, Tasmania, and the Australian Capital Territory decided on harmonising and extending daylight saving arrangements from April 2008.

The Western Australia DST was incorporated starting 3 December 2006.

New Zealand

The New Zealand Government announced its decision to extend the DST starting September — 2007. Clocks will go forward an hour a week earlier than usual - on the last Sunday in September and back an hour on the first Sunday in April. The Act administered by the Department of Internal Affairs is detailed at:

http://www.dia.govt.nz/diawebsite.nsf/wpg_URL/Services-Daylight-Saving-Index?OpenDocument.

Venezuela

The Venezuelan government formalized the intention to change Venezuela timezone to GMT-4:30 effective on 9 December 2007 at 3 AM local time. Previously, the GMT offset Venezuela had been following was -4:00. It is now changed to -4:30.

The announcement can be found at:

http://www.abn.info.ve/go_news5.php?articulo=112279&lee=18.

Argentina

The Argentine government formalized the intention to change Argentine timezone effective on 30 December 2007 at 0 AM local time until 16 March 2008. Previously, Argentina did not have any DST plans for 2007-08.

The announcement can be found at:

<http://www.telam.com.ar/vernota.php?tipo=N&idPub=87481&id=201230&sec=1&dis=1>.

There are similar updates to the DST of Canada, Bahamas, Bermuda, Brazil, and Uruguay in V5.1B-5.

BIND Updated to Version 9.2.8

This kit replaces the current version of BIND Version 9.2.5 with BIND Version 9.2.8. BIND 9.2.8 fixes the security issues that were faced in the BIND 9.2.5 version.

Updated Tru64 UNIX Documentation on docs.hp.com

Starting with Tru64 UNIX V5.1B-5, the updated documents associated with the Tru64 UNIX release will be posted on HP's Technical documentation website <http://docs.hp.com>. The updated documentation for this release includes the *HP Tru64 UNIX Release Notes for Version 5.1B-6*. and the *HP Tru64 UNIX and TruCluster Server Version 5.1B-6 Patch Summary and Release Notes* documents.

The updated documentation for the Internet Express for Tru64 UNIX and the Advanced Server for UNIX (ASU) will also be posted on <http://docs.hp.com>.

The existing documentation for Tru64 UNIX is provided on the Documentation V5.1B CD, included with the Tru64 UNIX media kit.

The manuals for the current release, manuals and documentation sets from previous releases of Tru64 UNIX, TruCluster software, ASU, Internet Express, and other products, are also provided online from the following Tru64 UNIX Documentation website:

<http://h30097.www3.hp.com/docs/>

Standards Conformance

Tru64 UNIX continues to conform to the UNIX 98 and POSIX standards. Several important commands and system calls have been updated to conform to the changes in standards, including `printf`, `pthread_mutexattr_getprotocol`, and `pwrite`.

For more information, see the specific release notes for these commands and calls in the *HP Tru64 UNIX and TruCluster Server Version 5.1B- Patch Summary and Release Notes*, and the `std_unix98` parameter of the `sys_attrs_generic` manpage.

New Keyword Added to `sshd2_config` Configuration File for `sshd` daemon

The new keyword, `AuthInteractiveFailureRandomTimeout`, adds a random delay to the existing `AuthInteractiveFailureTimeout` delay. For information on `AuthInteractiveFailureTimeout`, see the `sshd2_config` manpage .

The `AuthInteractiveFailureRandomTimeout` keyword can take a value from 0 to 100 (in seconds). The default is 2 seconds. To disable `AuthInteractiveFailureRandomTimeout`, specify a value of 0. When a non-zero value is specified for this keyword, a random number of milliseconds up to the number of seconds specified multiplied by 1000 is added to the server delay specified by `AuthInteractiveFailureTimeout`.

New Option to Ignore Processor Set Boundaries

A new command-line argument, `lockinfo`, has been added to solve issues related to processor-set boundaries. The new command takes `-ignore_pset` as an optional argument and when passed, enables the `lockinfo` command to ignore processor-set boundaries. However, it will honor the RAD (Resource Affinity Domain) set boundaries if the `-rad` option is used.

Support for Evaluating String Comparison Expressions as per POSIX Standards

The `sh-posix` built-in test is modified to evaluate string expressions as per the POSIX standard and can interpret "(" and "!" as operands in a string comparison operation.

To produce this POSIX compliant action, set the `STDS_FLAG` environment variable to ALL:

```
STDS_FLAG=ALL
```

If `STDS_FLAG` is not set or is set to NULL, the test function interprets "(" and "!" as operators in string comparison and reports wrong result. This was the default action before test was modified.

For example, consider a string comparison operation where "(" is passed as operand:

```
# test "(" = "abc"
```

The following message is displayed:

```
sh: test: Specify a parameter with this command.
```

This message indicates that the test function has failed to interpret "(" as an operand. With the flag set, "(" and "!" will be treated as valid operands.

`iconv` Converter Support Surrogate Pairs in Unicode

The `iconv` converter has been modified to fix the incorrect processing of surrogate pair characters in Unicode. In order to maintain compatibility, the new environment variable `ICONV_OLD_SURROGATE` is introduced. If this environment variable is set to a non-NULL value, `iconv` converter behaves in the same manner as before, that is, `iconv` converter continues to produce wrong results for Unicode surrogate pairs.

New Sysconfig Tunable to Reduce Contention on AdvFS Frag Files

A new sysconfig tunable `AdvfsFragGroupDealloc` has been introduced to set the frag group deallocation policy for the AdvFS filesystem. Using this tunable, you can enable or disable the frag group deallocation policy. The default is enabled.

File operations such as `rm` and `close`, which release a single frag, can trigger the frag group deallocation process when a list of free frags is encountered. This process holds a lock while processing the frag group. Any other process or thread that tries to manipulate the same frag group experiences a hang due to lock contention. The hang lasts for the duration of the frag group processing. This situation arises when the frag file of a fileset is large and too many files are present with frag. The `AdvfsFragGroupDealloc` tunable helps in disabling the frag group deallocation, which reduces the lock contention on the frag file.

The `AdvfsFragGroupDealloc` tunable can be added to the `/etc/sysconfigtab` file, and a value can be assigned as per the desired frag group deallocation policy. Placing the tunable in the `/etc/sysconfigtab` file will make the value persist across system reboots. Alternatively, `/sbin/sysconfig -r` can be used to assign the value for the tunable. However, this does not persist across system reboots.

On a cluster this tunable must be set on all the cluster members.

New rc.config Variables to Hide User Process Arguments and Environmental Variables for ps and w Commands

By default, the `ps` command displays a process's arguments and the `ps e` command displays a process's environmental variables. You can prevent users from viewing the arguments and environmental variables of other users' processes. To hide user process arguments and variables, enable the variables in the `/etc/rc.config.common` file:

```
# rcmgr -c set TBL_ARGUMENTS_DISABLE 1
# rcmgr -c set TBL_ENVIRONMENT_DISABLE 1
```

However, the root can always view the arguments and environmental variables of all users.

Similarly, the `w` command displays commands and their arguments. To prevent users from viewing commands and the arguments of other users' processes, enable the variable in the `/etc/rc.config.common` file:

```
# rcmgr -c set TBL_ARGUMENTS_DISABLE 1
```

However, the root can always view the arguments of all users.

Conformance to Open Group Standards

Set the `STDS_FLAG` environment variable to `ALL` so that `pthread_mutexattr_getprotocol()` conforms to the Open Group standard.

UNIX 98 Compliance with libc

Some `libc` functions from the `printf`, `scanf`, and `streams` family have been made to comply with UNIX 98 standards. These settings are enabled using the `sys_attrs_generic` variable, `std_unix98`. This variable (`std_unix98`) should not be set to the value of `STD_UNIX98_ALL` without consulting the Tru64 engineering team. For more information, see the `sys_attrs_generic` manpage.

Netstat Read Error on Structures in a Live System

When trying to read a structure, the `netstat()` command displays the following message:

```
netstat: read from /dev/kmem: No such device or address
```

This can result from `netstat` reading structures that are dynamically undergoing change on a live system. This is a transient problem that will be reported to the user.

O_APPEND Flag has no Effect on Behaviour of pwrite()

The `pwrite()` system call has been modified to conform to UNIX98 standard behavior. `O_APPEND` flag now will have no effect on the behaviour of `pwrite()`. The `sysconfig` tunable `pwrite_no_append` (in VFS subsystem) has to be set to 1 to enable this behavior.

smmsp User and Group Not Required for sendmail

The `smmsp` user, group, and the `/usr/var/spool/clientmqueue` directory were created as a future requirement for `sendmail` in the previous patch release v5.1B-4. Because, `sendmail` is not `smmsp` enabled, the `smmsp` user, group, and `/usr/var/spool/clientmqueue` will no longer be required. It is recommended that you remove these items if they are not being used for any other purpose on the system, including alternate `sendmail` implementation.

The following command displays how to delete the `clientmqueue` directory tree:

```
# rm -rf /usr/var/spool/clientmqueue
```

The following command displays how to delete the `smmsp` user and group:

```
# userdel smmsp
# groupdel smmsp
```



NOTE: Check the root directory and delete the `clientmqueue` directory, the user, and group related to the root directory for the patch kit install as follows:

```
#chroot $_ROOT /sbin/rm -rf /usr/var/spool/clientmqueue
#chroot $_ROOT /usr/sbin/userdel smmsp
#chroot $_ROOT /usr/sbin/groupdel smmsp
```

where `$_ROOT` is the alternate root directory

Possible Performance slowdown of Oracle 8.1.7 after Tru64 UNIX Rebranding

When kernel profiling and auditing were run on Oracle, under Version 5.1B-3, Asynchronous I/O + Direct I/O calls were seen. However, in Version 5.1B-4 and higher versions, no Asynchronous I/O + Direct I/O calls (other than AIO setup calls) were seen.

If this behaviour is seen on your system, you can modify `/etc/sysconfigtab` under `generic` to change:

```
version_banner = HP Tru64 UNIX
version_avendor = HP
version_vendor = Hewlett-Packard Company
```

to

```
version_banner = Compaq Tru64 UNIX
version_avendor = COMPAQ
version_vendor = Compaq Computer Corporation
```

Then reboot the system and check the Oracle performance.

Version 5.1B-5 or Higher Kit Requires Uninstallation of Internet Express System Authentication LDAP Module (IAELDAMXXX)

The Version 5.1 B-5 or higher patch kit installation fails if the Internet Express System Authentication LDAP Module (IAELDAMXXX) is installed on the system. To install the Version 5.1 B-5 or higher patch kit, perform the following steps:

1. Uninstall Internet Express System Authentication LDAP Module (IAELDAMXXX). For example,

```
setld -d IAELDAMXXX
```

where, XXX stands for the IAELDAM version.

2. Install the Version 5.1 B-5 or higher patch kit.
3. Install Internet Express System Authentication LDAP Module (IAELDAMXXX). For example:

```
setld -l IAELDAMXXX
```

where, XXX stands for the IAELDAM version.

IBM Tivoli Storage Manager (TSM) client problems fixed

The following issues with running the IBM Tivoli Storage Manager (TSM) are fixed in the current version:

- The TSM client performs full backups rather than incremental backups.
- The TSM client skips files, giving errors indicating the files were changed during the backup process, even when those files were not modified.

Enhancements Introduced in Prior Kits

The following sections describe some of the key features and enhancements that were first delivered in previous patch kits.

Enhanced Cluster Interconnect Extended to 100 KM

This release provides support for Enhanced Distance Clusters. An Enhanced Distance Cluster is a cluster in which the interconnect has been extended up to 100 km using a gigabit LAN Ethernet connection. An Enhanced Distance Cluster provides basic high availability services in the event of the loss of a single component. However, it does not include all of the high availability services provided by TruCluster Server. See [Appendix A \(page 191\)](#) for information about setting up and configuring an Enhanced Distance Cluster.

Cluster Cloning Offers Alternative to No-Roll Patching

This kit provides a new installation method, generically referred to as cloning, using a new tool named `dupclone`. The process consists of two primary steps:

- Creating an exact duplicate of an existing system on an alternate set of disk drives.
- Using `dupclone` to install the patch kit to the alternate disk set. After completion, the system can immediately be rebooted using the alternate disks.

See the *Patch Installation Instructions* document and the new `dupclone(8)` reference page for information about using `dupclone`. See “[dupclone Error Message Can Be Ignored](#)” (page 20) for information about a message you may see when using `dupclone`.

Link Aggregation Extended to Cluster LAN Interconnects

This kit provides enhanced support for link aggregation (LAG) by extending it to cluster LAN interconnects. It does this by decreasing the latency associated with time critical Interconnect Communication Subsystem (ICS) remote procedure calls and by increasing the available bandwidth of the LAN interconnect, thereby allowing increased interconnect throughput.

Latency is improved through multiple active interfaces decreasing queue sizes (link aggregation) and through the separation of ICS channels that prefer low latency from channels that require high bandwidth.

Throughput is improved from multiple active interfaces decreasing queue sizes.

The primary goal is to distribute cluster component channel traffic among the interfaces that form part of the link aggregation group.

These changes are implemented with the following three new `sysconfigtab` attributes, which configure the interfaces and create the LAG set:

`ics_tcp_lag0` Configures a LAN interface to be part of the LAG set.

ics_tcp_lag_dist	Specifies the lag traffic distribution algorithm for a LAG interface.
ics_tcp_lag_serv_weights	Based on this value, the channel traffic for an ICS channel is distributed over the LAG interface.



CAUTION: Do not modify the default value of these attributes unless instructed to do so by support personnel.

The following restrictions apply to the use of these attributes:

- Supported only on DEGPA (alt), DEGXA (bcm), and DE60x (ee) network interface cards (NICs).
- Supported only on Ethernet (802.3 CSMA/CD) links.
- NetRAIN virtual interfaces cannot be included in link aggregation groups.
- Ports must be operating in full duplex mode.
- Ports in the same link aggregation group must operate at the same data rate.
- Ports in a link aggregation group must be attached to the same system, either server-to-server or server-to-switch.

Link aggregation enables system administrators to combine two or more physical Ethernet Network Interface Cards (NICs) and create a single virtual link. Upper-layer software sees this link aggregation group as a single virtual interface for example: lag0.

The single virtual link can carry traffic at higher data rates than a single interface because the traffic is distributed across all of the physical ports that make up the link aggregation group.

For more information see the *Tru64UNIX Technical Overview* and the *Network Administration: Connections* manual. For tuning and configuration information see the *lag(7)*, *lagconfig(8)*, *sys_attrs_ee(5)*, *sys_attrs_lag(5)*, and *inet_local(4)* reference pages.

NetRAIN over LAG Supported.

With this release, it is now possible to run NetRAIN over link aggregation (LAG).

Previously, you could not simultaneously use NetRAIN for redundant network devices and LAG (trunking) on the same network cards. Although the use of LAG provided redundancy, it could not provide a redundant switch solution because all devices must be connected to the same switch.

With the installation of this kit, you can run NetRAIN over LAG; that is, have two or more LAG trunk groups contained within a NetRAIN set. Although only one LAG group will be active at one time, the benefit is that it allows the use of redundant switches with a high bandwidth LAG group.

To configure this support, first create your LAG groups, then place each group (lag0, lag1) into a newly configured NetRAIN set. Refer to the *lagconfig(8)* and *nr(7)* reference pages for details.

Support Provided for 2007 Changes to U.S. Daylight Savings Time

This kit updates `/etc/zoneinfo` time zone data files to incorporate the most recent changes in various time zones around the world, most notably the US Daylight Saving Time (DST) rule changes that were passed into law on August 8, 2005 and take effect in 2007.

That law moves the start of DST from the first Sunday of April to the second Sunday of March. It moves the return to Standard Time from the last Sunday of October to the first Sunday of November. These changes affect all US time zones and a number of other North American time zones in other countries as well.

BIND Updated to Version 9.2.5

This kit replaces the current version of BIND (V8.2.2) with BIND Version 9.2.5. (See the “Commands Must Be Run on BIND Systems After Kit Installation” (page 16) sections for

information about BIND actions to take when installing this kit.) This new version from the Internet Software Consortium represents a major rewrite of nearly all aspects of the underlying BIND architecture. Some of the important features of BIND 9 are:

- DNS Security
 - DNSSEC (signed zones)
 - TSIG (signed DNS requests)
- IP version 6
 - Answers DNS queries on IPv6 sockets
 - IPv6 resource records (AAAA)
- DNS Protocol Enhancements
 - IXFR, DDNS, Notify, EDNS0
 - Improved standards conformance
- Views
 - One server process can provide multiple views of the DNS namespace, for example, an “inside” view to certain clients, and an “outside” view to others.
- Multiprocessor Support
- Improved Portability Architecture

Library Calls for Fibre Channel HBA Added

This kit provide a wrapper library and HBA-specific library for the Emulex adapter, which conforms to the T11 FC-HBA (T11/1568-D Revision 14) specification.

The purpose of this specification is to provide a host bus adapter (HBA) programming interface for Fibre Channel management applications to gather information about devices in the network in a vendor-neutral way through a set of Application Programming Interface (APIs). As a result of this neutrality, applications do not depend on the platform they run on or for a specific HBA, and therefore will not need to be rewritten.

This kit supports the following APIs:

- HBA_GetVersion
- HBA_LoadLibrary
- HBA_FreeLibrary
- HBA_RegisterLibraryV2
- HBA_GetWrapperLibraryAttributes
- HBA_GetVendorLibraryAttributes
- HBA_GetNumberOfAdapters
- HBA_RefreshInformation
- HBA_RefreshAdapterConfiguration
- HBA_GetAdapterName
- HBA_OpenAdapter
- HBA_CloseAdapter
- HBA_GetAdapterAttributes
- HBA_GetAdapterPortAttributes
- HBA_GetDiscoveredPortAttributes
- HBA_GetPortAttributesByWWN
- HBA_GetPortStatistics
- HBA_GetBindingCapability
- HBA_GetBindingSupport
- HBA_SetBindingSupport
- HBA_GetFcpTargetMapping

- HBA_GetFcpTargetMappingV2
- HBA_GetFcpPersistentBinding
- HBA_SendScsiInquiry
- HBA_ScsiInquiryV2
- HBA_SendReportLUNs
- HBA_ScsiReportLunsV2
- HBA_SendReadCapacity
- HBA_ScsiReadCapacityV2

For information about each of these APIs, see the *Storage Management HBA API (SM-HBA)* standard (T11/1695-D) available at the www.t11.org website:

<http://www.t11.org/ftp/t11/pub/sm/hba/06-382v1.pdf>

The API information is listed under FC-HBA Function Calls.

New Cluster Command Sends ping Packets over TCP

A new TruCluster Server command, `clu_ping`, sends ping packages over the TCP layer rather than the Internode Communication Subsystem (ICS) layer on clusters with LAN as the interconnect. By default, ping packets are sent over the ICS layer.

For information about using this command, see the `clu_ping(8)` delivered in this kit.

sendmail Server Updated to Version 8.13.6

The `sendmail` server has been updated from Version 8.11.1 to Version 8.13.6. Key changes to the `sendmail` configuration file (`sendmail.cf`) include the following:

- The local mailer program, `bin/mail`, has been changed to `/usr/sbin/mail.local`
- The default database format has been changed from `dbm` to `btree`
- An additional security option has been added to the `imap deliver` program
- The IPC Mailer argument has been changed to `TCP`

The Version 8.13.6 `sendmail` server provides advanced features, including the following:

- Masquerading
- Virtual domain hosting
- Restricted relaying
- Milter functionality



NOTE: These features can be configured only with the `sendmail` provided with the HP Tru64 Internet Express Software distribution.

A new account, `smmsp`, is created as part of the `sendmail` installation process. This account is required for future enhancements of `sendmail`.

You can find information about `sendmail` Version 8.13.6 as follows:

- The `sendmail.org` website: <http://sendmail.org/>.
- The `sendmail` documentation provided with the HP Tru64 Internet Express Software distribution.
- The book *Sendmail* by Bryan Costales, and Eric Allman, published by O'Reilly & Associates, Inc.

The `sendmail` v8.13.6 reference pages were not updated in this release.

AdvFS Utilities Improved for Working with Metadata Files

The AdvFS vods utilities (`nvbmtpg`, `nvlogpg`, `nvtagpg`, `nvfragpg`, and `vsbmpg`) have been improved and enhanced to make them more useful when working with AdvFS metadata files.

New options have been added and some existing options have been improved. Review the revised reference pages included in this kit before using these enhanced tools.

mountd Daemon Gets New Port-Selection Option

A new option to the `mountd` daemon lets you specify a port number for `mountd` to bind to. Currently, when `mountd` starts it takes an arbitrary port number, which is different every time you boot your system. As a result, some applications may fail because the port number for the applications are defined in `/etc/services` and `mountd` may use one of them.

By using the new `mountd -p`, you can force `mountd` to bind to the specified port number instead of using the random port number. For example:

```
# mountd -p 1024
```

For more information, see the revised `mountd(8)` reference page that is installed with this kit.

envmond Daemon Modified to Use EVM Events

The `envmond` daemon has been modified to allow it to use EVM events instead of the `hwmgrr` command to determine the environmental status of the system. On systems with many sensors, this improvement may reduce or eliminate previously seen performance problems.

By default, `envmond` is configured to use the `hwmgrr` command (the poll method) for environmental monitoring. To configure `envmond` to use EVM events, set the `envconfig ENVMON_MODE` variable to `event` as follows:

```
# envconfig -c ENVMON_MODE=event
```

Because threshold values in event mode cannot be set to individual sensor, EVM events are generated only when existing hardware thresholds are exceeded. If you need to monitor individual sensors at thresholds different from the hardware thresholds, use the new `envconfig ENVMON_POLL_SENSORS` variable in conjunction with the `hwmgrr`. For example:

```
# envconfig -c ENVMON_POLL_SENSORS="58:59"
```

To then set the warning threshold for the sensor with ID 58, to 50.0 degrees Celsius, enter the following command:

```
# /sbin/hwmgrr -set attr -id 58 -a warning_threshold=500
```

To set the fault action for the sensor with ID 58 to `noshutdown`:

```
# /sbin/hwmgrr -set attr -id 58 -a fault_action=noshutdown
```

By specifying this set of commands, `envmond` uses the poll method for sensors 58 and 59 and the EVM event method for the rest of the sensors.

For more information, see the revised `envconfig(8)` reference page that is installed with this kit.

aha_chim Driver Problem Corrected

This kit corrects the following problems found in the `aha_chim` driver:

- The driver would fail to issue a "Bus Reset" instruction when a Bus Device Reset fails to complete for any reason other than a Bus Reset. With the installation of this kit, when a Bus Device Reset fails, a Bus Reset will be issued in order to clear and reset the target.
- The driver would fail to report the correct event identifier in the error entry of the binary error log when the error condition was caused by the target using an invalid tag ID, thereby resulting in an incorrect diagnosis of the problem. For example, an entry of the following type:

```
Event information: Adapter requested initialization,  
caller ID = 10
```

should have been:

```
Event information: Adapter requested initialization, caller ID =  
74.
```

Command Option Now Provides Additional EMX Driver Information

After installing this kit, issuing the following command for an EMX adapter will return the hardware revision, firmware revision, SAN address, and full duplex flag attributes:

```
# hwmgr -get att
```

New EMX Subsystem Attribute Turns on LLER for Tape I/O

This release provides a new attribute, `erp_ller`, to the EMX subsystem that allows you to turn on Link Level Error Recovery (LLER) for tape I/O. When enabled, the Emulex adapter attempts to successfully complete I/O that would have otherwise failed to a link error. If the adapter is unable to successfully complete the I/O, the I/O will be returned with an appropriate error.

Setting `erp_ller` to a value of 1 enables this feature. It is turned off by default due to issues seen with network storage routers and its handling of device resets. Command timeout errors may be returned if a device reset is issued when Link Level Error Recovery is enabled.

If you are experiencing failed tape I/O due to link issues, you can enable this feature and see if it helps.

To view the current setting of the attribute use the following command:

```
# sysconfig -q emx erp_ller
```

For more information, see the revised *emx(7)* reference page delivered in this kit.

Kernel Attributes Protect Against ICMP Security Vulnerability

A new kernel attribute delivered in this kit, `icmp_tcpseqcheck`, and an existing attribute, `icmp_rejectcodemask`, can protect your system against potential Internet Control Message Protocol (ICMP) security vulnerabilities. This release note describes these attributes and provides background information on the security issues. For information about setting these attributes, see the revised *sys_attrs_inet(5)* reference page delivered in this kit.

An overview of these attributes follows:

- `icmp_tcpseqcheck`

Mitigates ICMP attacks against the Transmission Control Protocol (TCP) by checking that the TCP sequence number contained in the payload of the ICMP error message is within the range of the data already sent but not yet acknowledged. An ICMP error message that does not pass this check is discarded. This behavior protects TCP against spoofed ICMP packets.

- `icmp_rejectcodemask`

A bitmask that designates the ICMP codes that the system should reject. The `icmp_rejectcodemask` attribute can be used to reject any ICMP packet type, or multiple masks can be combined to reject more than one type.

In the Requirements for Internet Protocol (IP) Version 4 Routers (RFC 1812), research suggests that the use of ICMP Source Quench packets is an ineffective (and unfair) antidote for congestion. HP therefore recommends using the `icmp_rejectcodemask` attribute to ignore ICMP Source Quench packets.

The ICMP type codes are in `/usr/include/netinet/ip_icmp.h`.

The ICMP (RFC 792) is used in the Internet Architecture to perform fault-isolation and recovery (RFC 816), which is the group of actions that hosts and routers take to determine if a network failure has occurred.

The industry standard TCP specification (RFC 793) has a vulnerability whereby ICMP packets can be used to perform a variety of attacks such as blind connection reset attacks and blind throughput-reduction attacks:

- Blind connection reset attacks can be triggered by an attacker sending forged ICMP "Destination Unreachable, host unreachable" packets or ICMP "Destination Unreachable, port unreachable" packets.
- Blind throughput-reduction attacks can be caused by an attacker sending a forged ICMP type 4 (Source Quench) packet.

Path MTU Discovery (RFC 1191) describes a technique for dynamically discovering the MTU (maximum transmission unit) of an arbitrary internet path. This protocol uses ICMP packets from the router to discover the MTU for a TCP connection path. An attacker can reduce the throughput of a TCP connection by sending forged ICMP packets (or their IPv6 counterpart) to the discovering host, causing an incorrect Path MTU setting.

caa_relocate Command Improved

The `caa_relocate -s source_member` command now allows the relocation of a specific resource from the `source_member`.

The command `caa_relocate -s source_member resource_name` will relocate the application resource `resource_name` only if it is running on the `source_member`. Otherwise it will return an error message.

See the revised `caa_relocate(8)` reference page delivered in this kit for more information.

collect Utility Improved in Several Ways

The `collect` utility has been enhanced to support a new `-c` option, which when specified instructs `collect` to gather local and remote I/O access statistics for disk and tape devices as seen by the Device Request Dispatcher (DRD) cluster subsystem in a TruCluster Server environment.

The `collect` utility has also been modified to enable it to support long device names.

The `collect(8)` reference page has been revised to reflect these changes.

Environment Variable Improves btcreate Kernel Build

This kit provides the means to allow the `btcreate` command to build the kernel with all options. Currently, if the kernel built with the current system configuration exceeds the firmware limit, `btcreate` will remove all options except DVDFS and CDFS. If the newly built kernel with CDFS and DVDFS also fails, `btcreate` then builds a kernel with mandatory options alone.

To build a kernel with all options, run `btcreate` by setting the following environment variable:

```
BTCREATE_MODE=VER-1-1
```

See the revised `btcreate(8)` reference page delivered in this kit for more information.

New Variable Aids Performance of AdvFS Administration Commands

A new `rc.config` variable, `ADVFS`, lets you control the boot time invocation of the `advfsd` daemon. This daemon is not necessary unless you are running the AdvFS graphical interface `dtadvfs`. Disabling `advfsd` from starting results in a better performance of AdvFS administration commands. See "Stopping Daemons May Speed Administration Performance" for more information about this problem.

The following list provides information on using the `ADVFS` variable to disable and enable the `advfsd` daemon on different types of systems:

- Run the following command to disable the `advfsd` daemon at boot time on a stand-alone system:


```
# /usr/sbin/rcmgr set ADVFS "no"
```
- Run the following command on any cluster member to disable the `advfsd` daemon at boot time on all members of a cluster:

- ```
/usr/sbin/rcmgr -c set ADVFSD "no"
```
- Run the following command to enable the `advfsd` daemon at boot time on a stand-alone system:

```
/usr/sbin/rcmgr delete ADVFSD
```
  - Run the following command on any cluster member to enable the `advfsd` daemon at boot time on all members of a cluster:

```
/usr/sbin/rcmgr -c delete ADVFSD
```

## New ftpd Command Option Prevents Login Delays

A new option to the File Transfer Protocol server daemon, (`ftpd -n`), can prevent login delays and time-outs in an environment where host name resolution is sluggish. It does this by disabling reverse lookups of remote host names.

This option is documented in the revised `ftpd(8)` included in this kit.

## New Features Added to kdbx Debugger

The `kdbx` command has been enhanced in several ways:

- A new cluster alias extension, `clua`, has been added to provide information about cluster aliases.
- New options, `-s` and `-v`, have been added to the `netstat` extension to expand its usefulness:
  - The `-s` option, when used alone, displays protocol statistics for all configured interfaces. When used with the `-i` option, `-s` displays interface statistics for all configured interfaces.
  - The `-v` option displays verbose information (including hardware addresses) about all interfaces that are configured on a system.
- A new option, `-p` has been added to the `inpcb` extension to display process ID (PID) information for each connection.

The revised `kdbx` reference page included in this kit describes the new `clua` cluster alias extension and the other new options.

## Modified rmvol Utility Allows Multiple Volume Removal

Modifications to the AdvFS `rmvol` utility now allow it to accept more than one volume for removal on the command line. In the following example, `rmvol` removes three volumes from a domain:

```
rmvol dsk5b dsk3a dsk4a rmvol_dmn1
rmvol: Removing 3 volume(s) from domain 'rmvol_dmn1'
rmvol: Removing volume '/dev/disk/dsk5b' from domain 'rmvol_dmn1'
rmvol: Removed volume '/dev/disk/dsk5b' from domain 'rmvol_dmn1'
rmvol: Removing volume '/dev/disk/dsk3a' from domain 'rmvol_dmn1'
rmvol: Removed volume '/dev/disk/dsk3a' from domain 'rmvol_dmn1'
rmvol: Removing volume '/dev/disk/dsk4a' from domain 'rmvol_dmn1'
rmvol: Removed volume '/dev/disk/dsk4a' from domain 'rmvol_dmn1'
rmvol: Removed 3 volume(s) from domain 'rmvol_dmn1'
```

Also, the new `rmvol -s` option performs a free-space check before beginning `rmvol` operations. If calculations determine that not enough free space will be available for the complete migration of all data for all volumes requested for removal, `rmvol` will fail before migrating any data. Upon failure, the amount of free space needed for complete migration of all data is displayed. For example:

```
rmvol -s dsk1a dsk3b test
rmvol: Removing 2 volume(s) from domain 'test'
rmvol: Not enough free space for complete migration of all volumes
requested for removal.
 Free space needed: 65592K
```

```
Free space available: 46296K
rmvol: Can't remove 2 volume(s) from domain 'test'
```

See the revised *rmvol(8)* reference page included in this kit for more information and additional examples.

## New disklabel Command Option Expands Partitions

A new option to the `disklabel` command lets you extend a partition that is currently in use. This option, `F`, is used with the `-e` option as follows:

```
/sbin/disklabel -e -F disk
```

For more information, see the revised *disklabel(8)* reference page included in this kit.

## Commands Modified to Conform to POSIX Standard

The following Tru64 UNIX commands have been modified to conform to the POSIX standard. For most of these commands, the modified action is initiated by using a new environment variable, `STDS_FLAG`.

- `awk`
- `cp`
- `ex`
- `chmod`
- `edit`
- `find`
- `rm`
- `uucp`
- `uudecode`
- `vi`

The following sections describe the changes to these commands.

### Changes to `ex`, `edit`, and `vi`

The `ex`, `edit`, and `vi` (`vedit/view`) commands have been modified so the POSIX compliant shell, `/usr/bin/posix/sh`, is the default shell when the `SHELL` environment variable is not set or is set to `NULL`.

Prior to this fix, `vi`, `ex` did not have a command line interpreter when the `SHELL` environment variable was set to `NULL`.

Setting `STDS_FLAG` to `ALL` produces the following POSIX compliant behavior:

If `C` or `S` is entered in command mode and more than part of a single line is affected, then `vi` saves the affected text in numeric buffers.

### Changes to `awk` and `nawk`

The `awk` and `nawk` commands have been modified to interpret numbers and the equal sign (=) as text strings when specified as arguments to "program text."

To produce this POSIX compliant action, set the new `STDS_FLAG` to `ALL`:

```
STDS_FLAG=ALL
```

When `STDS_FLAG` is set to `ALL`, variable names that do not begin with the alphabetic character or underscore are considered invalid.

If `STDS_FLAG` is not set or is set to `NULL`, `awk` interprets this use of numbers and the equal sign as numeric strings when specified as arguments to "program text." This was the default action before these commands were modified.

## Changes to chmod

The `chmod` command has been modified to force it to consider the `umask` when the `who` (`ugoa`) argument is not specified.

To produce this POSIX compliant-action, set the new `STDS_FLAG` to `ALL`:

```
STDS_FLAG=ALL
```

If `STDS_FLAG` is not set or is set to `NULL`, `chmod` does not consider the `umask` value while changing to the permissions specified. This was the default action before `chmod` was modified.

## Changes to cp

The `cp` command has been modified to enable compliance to the following POSIX requirements:

- When the `-i` and `-f` options are used together the `-f` should not disable a previous `-i` (that is, turn off prompting).
- When the `-f` is set and the target file cannot be opened for writing, `cp` unlinks the target file.

To produce this POSIX-compliant action, set the new `STDS_FLAG` to `ALL`:

```
STDS_FLAG=ALL
```

If `STDS_FLAG` is not set or is set to `NULL`, when the `-i` and `-f` options are used together the one specified last takes effect. This was the default action before `cp` was modified.

## Changes to ex

The `ex` command has been modified to return 1 as an exit status when a read-only option with write fails.

To produce this POSIX-compliant action, set the new `STDS_FLAG` to `ALL`:

```
STDS_FLAG=ALL
```

If `STDS_FLAG` is not set or is set to `NULL`, `ex` will return 0 as an exist status when a read-only option with write fails. This was the default action before `ex` was modified.

## Changes to find

The `find` command has been modified to not treat a hyphen (`--`) as special if it is first argument. Instead, it ignores the hyphen and lists the file containing the hyphen.

To produce this POSIX-compliant action, set the new `STDS_FLAG` to `ALL`:

```
STDS_FLAG=ALL
```

If `STDS_FLAG` is not set or is set to `NULL`, `find` will treat the first hyphen as special and exit with an error. This was the default action before `find` was modified.

## Changes to rm

The `rm` command has been modified to handle an excessive depth of files. Even if the pathname is longer than `PATH_MAX` by multiple times, `rm` will delete the directory with all its subdirectories and exit with value 0.

To produce this POSIX-compliant action, set the new `STDS_FLAG` to `ALL`:

```
STDS_FLAG=ALL
```

When `STDS_FLAG` is not set or set to `NULL`, `rm` will not delete files when the pathname exceeds `PATH_MAX` value. This was the default action before `rm` was modified.

## Changes to uucp

The `uucp` command has been modified so it can create a regular file when a directory with the same name already exists.

To produce this POSIX-compliant action, set the new `STDS_FLAG` to `ALL`:

```
STDS_FLAG=ALL
```

If `STDS_FLAG` is not set or is set to `NULL`, when `uucp`, attempts to create a regular file with the same name as an existing directory, the attempt fails and the file attributes are not changed. This was the default action before `uucp` was modified.

### Changes to `uudecode`

The `uudecode` command has been enhanced to recognize symbolic file mode.

For example, consider a case in which an editor was used to modify the first line of a source file of an encoded file from this:

```
begin 744 example.en
```

to this:

```
begin u=rwx,go=r example.en
```

The modified `uudecode` command would recognize the symbolic mode and create the file `example.en`.

To produce this POSIX-compliant action, set the new `STDS_FLAG` to `ALL`:

```
STDS_FLAG=ALL
```

If `STDS_FLAG` is not set or is set to `NULL`, when `uudecode`, will recognize only absolute file mode. This was the default action before `uudecode` was modified.

## New Generic Subsystem Attribute Corrects UNIX98 Standards Violations

A new tunable system attribute, `std_unix98`, has been added under the generic subsystem to cause the `waitpid()` and `poll()` system calls to conform to UNIX98 standard behavior.

See the revised `sys_attris_generic(5)` reference page delivered in this kit for more information. Refer to the `standards(5)` reference page for more information about industry standards and associated tags.

### `waitpid()` System Call

Prior to the installation of this kit, the `waitpid()` system call failed to conform to the following UNIX98 requirement:

A call to `pid_t waitpid(pid_t pid, int *stat_loc, int options)` when

- the calling process has `SA_NOCLDWAIT` set or has `SIGCHLD` set to `SIG_IGN` and
- has no unwaited for children that were transformed into zombie processes shall block until all of its children terminate, fail, and set `errno` to `ECHILD`.

The new `std_unix98` attribute enables `waitpid()` to conform to UNIX98 standard behavior.

For example, consider a situation in which a calling process has multiple children and no unwaited-for child zombie and you call `waitpid()` with a specific child PID:

- If you set `std_unix98=1` or `std_unix98=4`, `waitpid()` blocks until all of its children terminate (UNIX98 standard behavior).
- If you set `std_unix98=0` `waitpid()` blocks until any of its children exits.

### `poll()` System Call

Prior to the installation of this kit, the `poll()` system call fails to conform to the following UNIX98 requirement:

When no priority band has been written to on this `STREAM`, then a successful call to `int poll(struct pollfd fds[], nfd_t nfd, int timeout)` shall examine each element of the `fds` array for instances where the `POLLWRBAND` flag is set in the events member and data for a priority band greater than 0 can be written to the file descriptor specified by the `fd` member without blocking and

shall set the POLLWRBAND flag in the corresponding revents member when found.

When no writes have taken place on any of the priority bands, a call to `poll ( )` blocks will time out and return failure.

The `poll ( )` system call has been modified so if you need standards-compliant behavior, you can use the new `std_unix98` attribute.

If `std_unix98` is set to a value of either 1 or 2, then the `UNIX98_POLLWRBAND` bit (defined as a macro in the `/usr/sys/include/sys/param.h` file) gets set, which results in the internal processing becoming indifferent to whichever external mapping of `POLLWRBAND` is in play. By default, this bit in `std_unix98` is not set, so `poll ( )` will behave the same way as it does today.

See the revised *poll(2)* reference page delivered in this kit.

## New I/O Subsystem Attributes Can Improve Booting Speed

Three new I/O subsystem attributes control path registration during the boot process, allowing systems with multiple paths to a large number of devices to boot faster. The following list provides a brief description of these attributes. For additional information and settings, see the revised *sys\_attrs\_io(5)* reference page delivered in this kit.

- `boot_wait_hwc_reg`  
When disabled, causes a boot to the login prompt without waiting for hwc path registrations.
- `hwc_reg_cmplt_notify_type`  
Controls how you get notified when device registration is done.
- `hwc_registration_complete`  
Proves a query to determine if hwc path registration is complete.

By default, booting will wait for all hwc registrations to be completed. However, you can force the boot process to complete to the login prompt earlier by changing `boot_wait_hwc_regs` to 0. In either case `hwc_registration_complete` can be queried. This will be set to 1 as soon as registration is complete. In addition, you can also choose to receive a console message, an EVM event, or both when all paths have registered.

To get the biggest speed improvement when booting, you can elect to finish booting without waiting for path registration, which is not needed to access the storage subsystem (for example, an Oracle® database). However, if you do this, you temporarily have an incomplete hierarchy view from the commonly run `hwmgrr` command.

Each of the following actions can help you determine when to run the `hwmgrr` command in order to see the complete hierarchy:

- Enable the EVM notification option, log in, and start `evmwatch` to look for the EVM event, although the EVM event could have already occurred by the time you log in.
- Enable the console log message notification option and look in the messages file for the message.
- Query `sysconfig` to ensure the I/O `hwc_registration_complete` attribute is set before proceeding. This action can be used regardless of how you set the notification option.

## New Attributes Added to NFS and RPC Subsystems

Several new tunable attributes have been added to the NFS server subsystem, `nfs_server`, the NFS client subsystem, `nfs`, and the Remote Procedure Call (RPC) subsystem, `rpc`. Previously, the configurations produced by these attributes could only be changed by using the `dbx` command. Now, you can easily use and modify these kernel subsystem configurations with the `sysconfig` command.



The following list provides a brief description of these new attributes. For more information about setting the attributes, see the new *sys\_attrs\_nfs(5)* reference page delivered in this kit.

- `nfs_server`:
  - `nfs_write_gather` and `nfs3_write_gather`  
Improves NFS V2 and V3 performance by gathering several write requests, performing a single sync, and sending all of the replies.
  - `nfs_ufs_lbolt` and `nfs3_ufs_lbolt`  
Enables or disables a delay when NFS V2 and NFS V3 returns writes. This attribute affects write gathering for all file systems (not just UFS) for NFS V2 and V3.
- `nfs`:
  - `nfs_cto`  
Enables or disables Close-To-Open (CTO) consistency to reduce the number of client caches that provide applications with stale NFS data.
  - `nfs_quicker_attr`  
Enables or disables synchronous cache flush.
  - `nfs3_broken_lookup`  
Controls the frequency of console messages related to an NFS V3 problem on some servers where a file lookup would return erroneous data for the parent directory.
  - `do_client_readdirplus`  
Enables or disables the operating system from issuing the `readdirplus` procedure.
  - `nfs3_maxreadahead`  
Controls the number of outstanding read-aheads.
  - `nfs3_readaheads`  
Controls the number of read-aheads for NFS V3.
- `rpc`:
  - `use_fastsend`  
Enables or disables the optimization of client and server code used by NFS over UDP.
  - `use_fastroute`  
Enables or disables improved fastsend optimization that affect the NFS server.

## New cam Attribute Controls Path Usage

A new attribute, `cam_ccfg_aa_enable`, has been added to the `cam` subsystem to control preferred path usage.

When enabled, this attribute utilizes the target port group information from the storage controller to determine the optimal paths and use these optimal paths for I/O access. When disabled, all paths to that device are used.

Active-active asymmetric storage controllers may incur a performance penalty when accessed on non-optimal paths.

For more information, including setting this attribute, see the revised *sys\_attrs\_cam(5)* reference page delivered in this kit.

## LSM hot-sparing Improved

The Logical Storage Manager (LSM) command `volwatch` has been enhanced to improve LSM hot-sparing, which pro-actively replaces plexes that are based on failing storage devices and recovers their data.

Now when hot-sparing is performing a recovery, it will avoid using a plex that it is relocating — unless it has no other choice.

## New Option Changes Configuration File Used by `aliasd` Daemon

A new option, `custom_gated`, has been added to the `cluamgr` command to change the configuration file used by the `aliasd` daemon.

You can cause `aliasd` to use the file `/etc/gated.conf` instead of `/etc/gated.conf.memberX` as the `gated` configuration file and restart `gated` in either of the following ways:

- Specify the `cluamgr` command as follows:

```
cluamgr -r gated,custom_gated,start
```

- Modify the `/etc/rc.config.common` file by specifying the following command:

```
rcmgr -c set CLUAMGR_ROUTE_ARGS "gated,custom_gated"
```

After running the `rcmgr -c` command, restart network services by running the `rcinet` command on all cluster nodes:

```
rcinet restart
```

For more information about the `custom_gated` option, see the revised `cluamgr(8)` reference page included in this kit.

## `fsdb` Utility Now Operates on File System Image

The `fsdb` utility is now capable of operating on a file system image as well as a special file. The `name` argument will first be processed as a special file; should that fail, it will be processed as a regular file. To avoid conflict, an optional `f` argument will force the `name` argument to be processed only as a regular file. See the revised `fsdb(8)` reference page delivered in this kit.

## sendmail Log Problem Corrected

This kit corrects a problem with `sendmail` registration as a PSM (Process Set Manager) process. EVM would incorrectly log the following statements when `sendmail` was stopped or started or restarted:

```
PSM instance pid exited in category _unknown_ on node nodename
PSM instance pid created in category
unknown on node nodename
```

Where *pid* is the process ID of the `sendmail` daemon and *nodename* is the host where `sendmail` runs.

The `sendmail` program now correctly registers itself with the PSM and the same is reflected in log records.

## New Tunable Attribute Corrects NetRAIN Failover Problem

A new tunable attribute fixes a NetRAIN failover problem that occurs in a “quiet” network. In a two interface NetRAIN set, if the current active interface goes down, the secondary (backup) interface fails to become the new primary (active) interface. To correct this problem, set the new `nr_use_link_state` attribute as follows:

```
sysconfig -r netrain nr_use_link_state=1
```

For information about this attribute, see the revised `sys_attrs_netrain(5)` reference page delivered in this kit.

## New Attribute Controls Tape Driver Path Control

A new cam tape subsystem attribute, `enable_preferred_path`, lets you control preferred path behavior for a tape driver. Enabling this attribute (1) causes the tape driver to assign different paths to different tape drives. The default (0) disables preferred path behavior. The revised `sys_attr_cam(5)` reference page delivered in this kit describes the `enable_preferred_path` attribute.

A problem may exist when using preferred path behavior when you use a no-rewind tape device. If the application is not expecting the tape to change position between a previous tape close and a subsequent tape open, the data already on the tape may be lost with the next write command, possibly resulting in unusable backups. This problem only occurs when a system has multiple initiators.

The following steps will prevent this problem from occurring:

- Before performing a backup, reserve the device and lock down the path as follows:  

```
mt -f /dev/ntape/tape0 reserve
```
- After performing the backup, release the device and unlock the path as follows:  

```
mt -f /dev/ntape/tape0 release
```

## pr Command Behavior Now Works as Described in Reference Page

The `pr` command has been modified to handle the `i [character] [gap]` option (which replaces multiple space characters with tab characters) so that it performs as documented in the `pr(1)` reference page.

## Kit's Session Log Made More Useful

If you view the session log for this release, you may notice that it is smaller than it has been in the past. We edited this file to remove non-essential information that the system generates automatically.

## Sys\_Check Version 143 Provided

This kit includes `Sys_Check` Version 143. However, HP recommends that you visit the `sys_check` website to download and install the version there if it is more recent than Version 143:

[http://h30097.www3.hp.com/sys\\_check/](http://h30097.www3.hp.com/sys_check/)

## New Variables Protect Against Attack

This kit provides two new kernel tunable variables, `tcp_rst_win` and `tcp_syn_win` to protect systems against potential vulnerabilities called TCP RST attack and TCP SYN attack. For more information, see “Potential Security Vulnerability Identified” and a revised `sys_attr_inet(8)` reference page, which is installed with this kit.

## New fuser Option Aids Query Search

A new option, `-a`, has been added to the `fuser` command to expand a query to search of all cluster members. See the `fuser(8)` reference page for more information.

## New /etc/printcap Option Provided

This kit provides a new boolean `/etc/printcap` option, `sr`, to suppress the reprinting of jobs under conditions that indicate to the print daemon that a reprint is needed. The syntax for this entry is similar to that of the `sh` (suppress header) option.

You can use this option to suppress an unexpected or unneeded reprinting of jobs that are completed but are reprinted a second time due to miscommunication between the printer and the print daemon.

Be aware that if you set this option, incomplete jobs that trigger reprint conditions will not reprint.

A fix to remote job reprinting that this patch kit provides can trigger reprints which, under conditions previously described, do not appear to be needed.

## Support for the Name Services Switch Added

The Name Service Switch (NSS) has been added to Tru64 UNIX as a replacement for the `svc . conf` database service selection. The NSS provides a more extensible database service selector and supports a dynamic list of databases. Using the NSS allows you to add LDAP as a source for `netgroup` data.

Configuring the NSS converts entries from the `/etc/svc . conf` file into entries for `/etc/nsswitch . conf` file. The `/etc/svc . conf` is then only used for pre-nsswitch statically-built applications and `sendmail`. For more information about this feature, see *nsssetup(8)*, *nsswitch(4)*, and *nss2svc(8)*

## New Hardware Support

This patch kit provides the following new hardware support.

### Support for 64 Processor AlphaServer GS1280 Systems

This patch kit provides support for AlphaServer GS1280 systems configured with 64 processors.

### Support for AlphaServer and AlphaStation DS15 Systems

The AlphaServer/AlphaStation DS15 3U Systems include:

- Alpha 1 GHz CPU with 2 MB onboard ECC cache
- 512-MB, 1 GB, or 2 GB SDRAM memory - expandable to 4 GB
- Onboard dual 10/100 BaseT Ethernet ports
- Four 64-bit PCI expansion slots
- Onboard Ultra160 SCSI controller

### HP StorageWorks FCA2384

Support has been added for the FCA2384 - 2 GB, 64-Bit/133 MHz PCI- X-to-Fibre Channel Host Bus Adapter.

## Production Version of Motif 2.1 Provided

This kit replaces the Motif 2.1 Advanced Developer's Kit (ADK) with a production version of Motif 2.1. This new version will be supported in future Version 5.1B releases. The production version of Motif 2.1 will also be available for downloading from the Web.

## Protection Against Buffer Overflow Exploitation Added

This kit provides a 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.

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

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

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

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

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

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

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



---

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

---

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

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

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

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

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

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

See `chatr(1)` for information about the `chatr` utility.

## Enhancements to pmgrd Daemon and collect Utility

Patches in this kit provide enhancements to the performance manager metrics server daemon, `pmgrd`, and the `collect` utility.

### Performance Manager Metrics Server Daemon (`pmgrd`)

The following features have been added to `pmgrd`:

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

- Support for monitoring the AdvFS statistics.

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

- The domain name
- The fileset name
- Number of files and blocks
- Soft and hard limits of files
- Soft and hard limits of blocks
- The status of user and group quotas
- Grace time and fileset clone information

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

The `collect` utility displays these new AdvFS statistics. (See “New Features Added to `collect` Utility”).

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

```
sysman advfs
```

See `pmgrd(8)` for more details.

### New Features Added to `collect` Utility

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

- AdvFS monitoring capability. (See “Enhancements to `pmgrd` Daemon and `collect` Utility” for a list of AdvFS metrics that are monitored.)  
Enables `collect` to report AdvFS volume I/O queue and fileset vnode operation statistics. You can specify the domain or fileset to be monitored, using the `-A` option.
- Viewing CPU and memory metrics on a per Resource Affinity Domain (RAD) basis.  
When run on a NUMA platform, enables `collect` to automatically retrieve CPU and memory metrics for each RAD.

See `collect(8)` for more details.

## File System Management Applications Enhanced

Enhancements to the SysMan Menu file system management applications delivered in this patch kit significantly improve their performance.

## New Control Option for /usr/sbin/audit\_tool Command

This kit provides the following new control option for the /usr/sbin/audit\_tool command:

```
/usr/sbin/audit_tool -. [path]
```

This option causes audit\_tool to use [path] for the archive/recovery directory that contains archived audit logs, thereby overriding the directory specified in the audit log, which by default is /var/audit.

## Change to envmond Improves Performance on Some Systems

This patch kit modifies the Environmental Monitoring daemon, envmond, to improve performance on systems with many sensors.

With the default monitor period of envmond (ENVMON\_MONITOR\_PERIOD), systems having large number of sensors may experience a performance degradation. The changes to envmond address this problem by polling sensors at a lower frequency.

## Release Notes Introduced in Prior Kits

The release notes in this section were included in previous patch kits.

## Authentication Choice Affects sftp Transfer Rate

The performance of secure FTP (sftp) will be always less than ftp due to the authentication and encryption involved in sftp communication. The transfer rate of sftp depends on the type of authentication it employs. You can achieve better transfer rate performance by choosing the Message Authentication Code (MAC) algorithm hmac-md5 for authentication, but at the cost of security. The default MAC is hmac-sha1, which is more secure than hmac-md5. See sftp(1) for information about secure FTP and ssh2\_config(4) for information about supported MACs and ciphers.

## Tru64 UNIX Rebranding Results in File Changes

As a consequence of the rebranding of Tru64 UNIX from the Compaq name to HP, the following files have changed:

- version.abbrev\_vendor from COMPAQ to HP
- version.banner from Compaq Tru64 UNIX to HP Tru64 UNIX
- version.vendor from Compaq Computer Corporation to Hewlett-Packard Company

The .mrg..sysconfigtab file has been modified to incorporate these changes into the generic sysconfig subsystem in the /etc/sysconfigtab file.

If the rebranding of HP Tru64 UNIX version information impacts any applications or layered products, you can manually change generic system version attributes. See the sysconfigtab(4) and sys\_attrs\_generic(5) reference pages for more information on how to modify generic system version attributes.

## Insight Manager Components DUMP Core

Some Insight Manger components included in Tru64 UNIX Version 5.1B-4, such as cpq\_mibs and the config\_hmmod, and sysman\_hmmod daemons, may core dump during reboots.

You can correct this problem by installing the latest version of the Insight Manager. At the time Version 5.1B-4 was released, the most current Insight Management Agents kit was Version 3.6. You can download this version from the HP Insight Management Agents for Tru64 UNIX website:

<http://h30097.www3.hp.com/cma/>

If you have not installed Version 3.7 and the Insight Manager processes do not run after rebooting your system, restart them using the Insight Manger startup scripts.

## Autoloader Firmware Upgrade Changes WWND

A firmware upgrade to v1.50 or N14r on the 1x8 Autoloader causes the WWID to change. As a result, the existing device associated with the media changer is no longer accessible. For complete details see the Customer Advisory available at:

<http://h30097.www3.hp.com/unix/erp/c00753663.html>

## Some Smart Array Errors May Not Be Recoverable

When booting your system you may see a message similar to the following:

```
Smart Array at ciss(1) not responding - disabled.
```

A system reboot may be able to re-enable the hardware. If that does not work, you need to call Field Service and have the unit repaired.

## Do Not Use dxarchiver to Verify Bootable Tape

Do not use `dxarchiver` command to verify a bootable tape. Instead, use the `mt` and `restore` commands as follows:

```
mt fsf 1
restore -i <device>
```

The first command skips the first file on the tape.

When preparing for a `btcreate` session, verify the size of the file system to ensure that you have sufficient tape volumes, depending on the maximum storage capacity of your tape device. The `btcreate` command prompts you to load a new tape volume if it runs out of storage space. Label the tapes in sequence.

## securenets File Requires localhost Entry

If the `/var/yp/securenets` file is in use as part of NIS, it must contain the following localhost entry:

```
255.255.255.255 127.0.0.1
```

If the `/var/yp/securenets` file is used without a localhost entry you will see severe delays on logins. See `ypserv(8)` for more information.

## SIA sialog Use Limitation Required

The Security Integration Architecture (SIA) `sialog` logging process is only intended for use in debugging SIA problems. It should not be enabled for extended periods of time. Doing so can cause login delays or other problems.

Use the `audit` subsystem to monitor authentications on the system, not the `sialog` process

To disable `sialog` debug logging, delete the `/var/adm/sialog` file. For more information, see the `sialog(4)` and `sia_log(3)` reference pages and the Tru64 UNIX *Security Programming* manual.

Note that when used in a TruCluster Server cluster, the `sialog` file is a cluster-wide file.

## Change to executable\_data Attribute Requires Running Script

Prior to setting the tunable attribute `executable_data` to a non-zero value, you must run the following script:

```
/usr/sbin/javaexecutedata
```

## Potential Security Vulnerability Identified

The industry standard TCP specification, RFC793, has a vulnerability in which an attacker can reset established TCP connections using the TCP RST (Reset) or SYN (Synchronize) flags.

These packets need to have source and destination IP addresses that match the established connection as well as the same source and destination TCP ports.



The fact that TCP sessions can be reset by sending suitable RST and SYN packets is a design feature of TCP. According to RFC 793, an RST or SYN attack is only possible when the source IP address and TCP port can be forged (also called spoofed). In that case TCP sessions, including Telnet, SSH, SFTP and HTTP may be disconnected without warning. TCP sessions that have been disconnected can be re-established.

Normally, a TCP SYN packet (request for a new connection) that arrives on a server using a matching IP address, port number, and matching sequence number for an existing connection causes a TCP RST packet to be returned to the client. An attacker can guess the proper sequence number, along with the port and IP addresses, to cause an existing connection to be terminated with a TCP RST.

When a client is rebooted without closing an old connection to the server, a subsequent attempt to connect to the server that matches the old connection tuple and sequence number will require a TCP RST in order to purge the old (stale) connection.

HP has addressed these potential vulnerabilities, called TCP RST attack and TCP SYN attack, by providing two new kernel tunable variables, `tcp_rst_win` (TCP RST window) and `tcp_syn_win` (TCP SYN window).

These variables mitigate the TCP reset attack by reducing the window sizes in which a TCP RST/SYN packet will be accepted by the Tru64 UNIX system.

The attributes for these variables are described in a revised `sys_attrs_inet(5)` reference page included in this kit.

After the patch kit is installed, you can adjust the variables using the `sysconfig` and `sysconfigdb` commands, as described in the following sections.

#### Adjusting the `tcp_rst_win` Variable

You can adjust the TCP RST window variable, `tcp_rst_win`, as follows:

```
sysconfig -q inet tcp_rst_win
inet:
 tcp_rst_win = -1

sysconfig -r inet tcp_rst_win=2048
tcp_rst_win: reconfigured

sysconfig -q inet tcp_rst_win
inet:
 tcp_rst_win = 2048

sysconfig -q inet tcp_rst_win /tmp/tcp_rst_win_merge

sysconfigdb -m -f /tmp/tcp_rst_win_merge inet

sysconfigdb -l inet
inet:
 tcp_rst_win = 2048
```

#### Adjusting the `tcp_syn_win` Variable

You can adjust the TCP SYN window variable, `tcp_syn_win`, as follows:

```
sysconfig -q inet tcp_syn_win
inet:
 tcp_syn_win = -1

sysconfig -r inet tcp_syn_win=2048
tcp_syn_win: reconfigured

sysconfig -q inet tcp_syn_win
inet:
 tcp_syn_win = 2048
```

```
sysconfig -q inet tcp_syn_win /tmp/tcp_syn_win_merge
sysconfigdb -m -f /tmp/tcp_syn_win_merge inet
sysconfigdb -l inet
inet:
tcp_syn_win = 2048
```

## Modification to Changer Driver May Affect Some Applications

As a side effect of resolving issues with multiple access to the changer, the changer driver now requires a short period of exclusive access to the changer device as part of opening the device. For applications that have several threads or processes accessing a single changer simultaneously, this can result in waits for access to the changer device in the process of an open call. That wait can be lengthy as some changer commands can have long response times.

In general this behavioral change will not affect the overall throughput to a changer device, as this wait would have occurred at the time of any I/O (for example, IOCTLs) to the changer.

If having the changer wait in this fashion presents a problem, the old behavior can be approximated by passing either the `O_NONBLOCK` or `O_NDELAY` flags at the open of the changer device. In that situation the first actual I/O (usually an IOCTL) may incur the wait as the open is partially delayed in that case.

## Data Sorting of Audit Records May Be Required on Single CPU System

The `net_tcp_stray_packet`, `net_udp_stray_packet`, and `net_tcp_rejectd_conn` network events are handled by the audit subsystem differently from other auditable events. As a result, these events may be placed into the audit log out of order with respect to other events.

Previously, the sorting of audit data on single CPU systems was unnecessary. This changed, however, when the capability for auditing these network events was introduced. Now, to view these network events in order with respect to other events, you must sort the data on a single CPU system. To do this, use the `audit_tool -S` command.

## new\_wire\_method Tunable Attribute Retired

The tunable attribute `new_wire_method` has been retired. After you install this kit, setting `new_wire_method` to either 0 or 1 will no longer affect your system.

## Stopping Daemons May Speed Administration Performance

When using AdvFS administration commands, the `advfsd` and `smsd` daemons rescan filesets, domains and volumes for system information. Depending on the number of filesets, domains, and volumes, you may experience a pause — sometimes quite long — between the commands.

If you experience this performance degradation, you may want to stop `advfsd` (required for `dtadvfs`, the AdvFS graphical user interface) and `smsd` (required for SysMan Station) daemons before running multiple AdvFS administration commands.

See “New Variable Aids Performance of AdvFS Administration Commands” for information on disabling the `advfsd` daemon at boot time.

To temporarily stop the daemons enter the following commands:

```
/sbin/init.d/advfsd stop
/sbin/init.d/smsd stop
```

To restart the daemons enter the following commands:

```
/sbin/init.d/advfsd start
/sbin/init.d/smsd start
```

## sendmail Application Size/Length Limits Can Cause Problems

When upgrading older releases of `sendmail`, be aware that the 5.1B version of `sendmail` has MIME header/content marker size limits and message header length limits. These limits have been added to stop a Denial of Service (DoS) attack on the `sendmail` server. The values default to the following:

```
MIME Header Length Size = 2048 characters
MIME Content Marker Size = 1024 characters
```

The `MaxHeadersLength` value is the maximum message header length allowed and its size can be installation dependent (the value defaults to 8192 bytes).

Some legacy applications may be affected by this security addition if the application is sending mail messages with long lines of text and no new-line markers. These limitations may cause `sendmail` to insert a carriage return at these boundaries.

To revert back to the old `sendmail` behavior, do the following:

1. Verify the `V2/Digital` header line is in the `/var/adm/sendmail/sendmail.cf` file. If the line is there, proceed to step 2. If it's not there, add it above the `# predefined` line. For example:

```
vi sendmail.cf
```

```
#####
V2/Digital
```

```
predefined
```

2. Add the following lines to the `/var/adm/sendmail/sendmail.cf` file:

```
O MaxMimeHeaderLength=0/0
O MaxHeadersLength=-1/-1
```

3. Restart `sendmail`

## Increasing RDG max\_objs Value Recommended

For certain applications where Oracle instances are running in a cluster and Memory Channel is used as the interconnect, console messages of “`rdg: out of objects`” may occur.

Tuning the `sysconfigtab` value `max_objs` (under the `rdg` subsystem) can eliminate these messages. We recommend doubling your current value.

Because this parameter is not dynamic, you can only change it by modifying the `sysconfigtab` file and rebooting your system. After doing this, observe your cluster to see if the messages have been eliminated.

You can set this value to a maximum of 50,000.

## Reboot May Resolve Problem with Smart Array Controller

If a problem with your Smart Array controller generates the following message, try rebooting your system:

```
Smart Array at ciss(1) not responding - disabled.
```

If the reboot does not re-enable the hardware, you will need to call your HP support representative to have the unit repaired.

## Additional Steps for IPsec Connections

This kit fixes a potential security vulnerability in IP security (IPsec). If you have one or more IPsec connections configured on your system, you need to ensure that you have restricted access to each IPsec connection based on the identity of the remote hosts. You can accomplish this after installing this kit by starting the IPsec SysMan configuration tool from the command line:

```
sysman ipsec
```

Once you have started SysMan, you will need to modify the configuration of each IPsec and IKE connection to add the identity of the remote hosts that are allowed to connect.

You enter this information on the third dialog box you see during the connection configuration wizard; the dialog box is titled "Manage IPsec: Add/Modify Connection: IKE Proposal." Although you can leave the "Restrict To The Following Remote IDs" list empty, doing so will mean that any identity given to the local machine by the remote hosts will be considered valid as long as they send the correct certificate or preshared key.

## Potential NFS Duplicate Request Cache Scalability Limitation with High Loads and Uncharacteristic File Access Behavior on Clustered NFS Servers

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 indistinguishable 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 "Tuning the NFS Server Duplicate Request Cache".

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
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 |

## Tuning the NFS Server Duplicate Request Cache

The NFS server maintains a list of recently completed non-repeatable 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 re truncation of a file.

This patch kit provides a tuning variable, `nfs_dupcache_size`, to control the size of the NFS server's duplicate request cache, which 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.

## Performance of hwmgr Commands on Large System Configurations

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.
- On a maximally configured AlphaServer GS1280 system with highly interconnected storage, the `hwmgr -view hierarchy` command generates thousands of lines of output.

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.

## LSM Spin Lock Issue

A patch in this kit addresses a spin lock issue in the LSM kernel that may occur under extremely heavy I/O loads on multiprocessor systems.

To reduce the need for certain spin locks in the kernel I/O code, you can set a new `sysconfigtab` variable, `Max_LSM_IO_PERFORMANCE`, to 1 (the default is 0). Doing this will increase LSM I/O performance if it is found that performance is degraded because of a highly contentious spin lock.

Note that using this spin lock performance feature reduces some of the debugging statistics that are normally maintained.

In order to use this feature, you must allow at least one LSM I/O daemon (`voliod`). The `voliod` daemon was changed to prevent the number of LSM kernel I/O daemons from being set to zero if this spin lock performance feature is turned on.

The change to the following `voliod` command produces an error and the number of LSM kernel I/O daemons remain unchanged:

```
voliod -f set 0
lsm:voliod: ERROR: VOL_IO_DAEMON_SET failed: Permission denied
```

## Possible Problem when Processing Many Command Parameters

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

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

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

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

You can also use the `xargs` command to break a long argument list into smaller lists. For more information, see the `xargs(1)` reference page.

## Loading Firmware from a BOOTP Server

The `fwupgrade` command has been modified to allow the specified firmware update image to be loaded from a BOOTP server in a connected network. This process must use the `bootpd` daemon.

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

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

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

```
Example bootptab file for BOOTP support
```

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

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

```
#
```

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

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

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

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

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

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

## Changes to tar, pax, and cpio Behavior

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/
```

## Changes to vdump and vrestore Disallow Larger Record Sizes

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 the installation of this kit, the valid range is 2 through 64 KB blocks.

## Problem Seen on Systems with Smart Array Controller

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

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

If this occurs, take the following steps:

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

```
ciss:
ciss_throttle_threshold=5
```

2. Execute the following command:

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

3. Reboot your system:

```
shutdown -r now
```

## Broken Links Reported During Baselining

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

```
Phase 4 - Report changed systemfiles and missing files
=====
This phase provides information to help you make choices later in
this process. It reports both 'missing' and files whose origin
cannot be determined. Some of these files may affect patch
installation. You will want to consider this information when you
later make decisions in phase 5.

* list of changed files with unknown origin:

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

Press RETURN to proceed...

You can disregard this information. The presence of these broken links will not affect your system operation, the installation of `dupatch` or `dupatch` tools, the successful installation of patches, or the rebuilding of kernels on the system.

## Russian Keyboard

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

```
// KEY <AD09 can be replaced by an extra keycap.
// If you replace it with the extra keycap, please uncomment
// the following definition and comment out the original one.
//
```



```

// key <AD09 {
// symbols [Group1]=3D [o, o],
// symbols [Group2]=3D [Ukrainian_i, Ukrainian_I]
// };
 key <AD09 {
 symbols [Group1]=3D [o, o],
 symbols [Group2]=3D [Cyrillic_shcha, Cyrillic_SHCHA]
 };

```

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

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

### Time Loss on Systems with Firmware Lower Than V6.4-12

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

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

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

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

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

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

As a workaround to the problem, you can modify one of the following two files and then reboot your system for the new setting to take effect:

- `/etc/rc.config`

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

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

- `/etc/sysconfigtab`

At the end of the file, add the following line:

```
marvel_srvmgmt: MV_Env_Support = 0
```

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

### CPU Offline Restrictions

The Primary CPU cannot be taken off line.

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

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

### Problem with Capacity-on-Demand Process

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

With the capacity-on-demand process, the `codconfig [cpu_id_list]` command lets you specify which CPUs you have paid for and which are spares. The command is supposed to mark the others as spare and then take them off line. Once a CPU is marked as spare, the `hwmgr`

command and Manage CPUs suitlet will not let you put them on line until you use the `ccod -l` or `ccod -p` command to either loan or purchase the CPU.

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

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

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

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

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

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

### Repeated Reboots May Cause Panic

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

### Caution on Updating to Version 5.1B with DEGXA NICs

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

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

## Changes to Reference Pages

There are no new or updated reference pages in this kit. Those in Table 2-1 "Reference Pages Revised in Previous Version 5.1B Patch Kits" were included in earlier Version 5.1B patch kits and are installed with this kit if they were not previously installed.

**Table 2-1 Reference Pages Revised in Previous Version 5.1B Patch Kits**

|                         |                             |                          |                           |
|-------------------------|-----------------------------|--------------------------|---------------------------|
| <i>advfsd(8)</i>        | <i>aliases(4)</i>           | <i>disklabel(8)</i>      | <i>dsfmgr(8)</i>          |
| <i>dumpfs(8)</i>        | <i>fctl(2)</i>              | <i>find(1)</i>           | <i>fread(3)</i>           |
| <i>freezefs(2)</i>      | <i>gethostbyaddr(3)</i>     | <i>gethostbyname(3)</i>  | <i>getlogin(2)</i>        |
| <i>getrlimit(2)</i>     | <i>iconv_intro(5)</i>       | <i>ksh(1)</i>            | <i>lockinfo(8)</i>        |
| <i>mail.local(8)</i>    | <i>mailq(1)</i>             | <i>mailstats(8)</i>      | <i>makemap(8)</i>         |
| <i>mlock(3)</i>         | <i>mount(2)</i>             | <i>mount(8)</i>          | <i>netstat(1)</i>         |
| <i>newaliases(1)</i>    | <i>praliases(1)</i>         | <i>proplist(4)</i>       | <i>ps(1)</i>              |
| <i>sendmail(8)</i>      | <i>sftp2(1)</i>             | <i>sialog(4)</i>         | <i>signal(2)</i>          |
| <i>ssh2(1)</i>          | <i>ssh-agent2(1)</i>        | <i>sshd2_config(4)</i>   | <i>sys_attrs_advfs(5)</i> |
| <i>sys_attrs_cam(5)</i> | <i>sys_attrs_generic(5)</i> | <i>sys_attrs_inet(5)</i> | <i>sys_attrs_vfs(5)</i>   |
| <i>sys_attrs_vm(5)</i>  | <i>unmount(2)</i>           | <i>unmount(8)</i>        | <i>useradd(8)</i>         |
| <i>vacation(1)</i>      | <i>vfast(8)</i>             | <i>vipw(1)</i>           | <i>w(1)</i>               |
| <i>aio_return(3)</i>    | <i>audt_tool(8)</i>         | <i>awk(1)</i>            | <i>binlogd(8)</i>         |
| <i>btcreate(8)</i>      | <i>btextract(8)</i>         | <i>caa_relocate(8)</i>   | <i>chatr(1)</i>           |

**Table 2-1 Reference Pages Revised in Previous Version 5.1B Patch Kits (continued)**

|                              |                              |                              |                              |
|------------------------------|------------------------------|------------------------------|------------------------------|
| <i>chmod</i> (1)             | <i>clu_ping</i> (8)          | <i>clua_services</i> (4)     | <i>cluamgr</i> (8)           |
| <i>codconfig</i> (8)         | <i>collect</i> (8)           | <i>cp</i> (1)                | <i>csch</i> (1)              |
| <i>dd</i> (1)                | <i>ddr.dbase</i> (4)         | <i>dig</i> (1)               | <i>disklabel</i> (8)         |
| <i>dnssec-keygen</i> (8)     | <i>dnssec-makekeyset</i> (8) | <i>dnssec-signkey</i> (8)    | <i>dnssec-signzone</i> (8)   |
| <i>dump</i> (8)              | <i>dupclone</i> (8)          | <i>dxshutdown</i> (8)        | <i>edauth</i> (8)            |
| <i>emx</i> (7)               | <i>envconfig</i> (8)         | <i>envmond</i> (8)           | <i>EvmEventPost</i> (3)      |
| <i>evminfo</i> (1)           | <i>EvmVarSet</i> (3)         | <i>ex</i> (1)                | <i>find</i> (1)              |
| <i>freezefs</i> (8)          | <i>fsdb</i> (8)              | <i>ftpd</i> (8)              | <i>fuser</i> (8)             |
| <i>fwupgrade</i> (8)         | <i>getaddrinfo</i> (3)       | <i>gethostbyaddr</i> (3)     | <i>getnameinfo</i> (3)       |
| <i>host</i> (1)              | <i>hwmgr_show</i> (8)        | <i>ifconfig</i> (8)          | <i>ip</i> (7)                |
| <i>ip6rtrd</i> (8)           | <i>ip6rtrd.conf</i> (4)      | <i>javaexecutedata</i> (8)   | <i>kdbx</i> (3)              |
| <i>kdbx</i> (8)              | <i>ksh</i> (1)               | <i>lan_config</i> (8)        | <i>ldapcd.conf</i> (4)       |
| <i>ldapusers.deny</i> (4)    | <i>rndc</i> (8)              | <i>mdc.conf</i> (5)          | <i>rndc-confgen</i> (8)      |
| <i>migrate</i> (8)           | <i>mountd</i> (8)            | <i>mt</i> (1)                | <i>named</i> (8)             |
| <i>named.conf</i> (5)        | <i>named-checkconf</i> (8)   | <i>named-checkzone</i> (8)   | <i>netgroup</i> (4)          |
| <i>netstat</i> (1)           | <i>newfs</i> (8)             | <i>nvbmtpg</i> (8)           | <i>niffconfig</i> (8)        |
| <i>niffmt</i> (7)            | <i>nsdispatch</i> (3)        | <i>nslookup</i> (8)          | <i>nss2svc</i> (8)           |
| <i>nssetup</i> (8)           | <i>nsswitch.conf</i> (5)     | <i>nsupdate</i> (8)          | <i>ntp.conf</i> (4)          |
| <i>nvbmtpg</i> (8)           | <i>nvfragpg</i> (8)          | <i>nologpg</i> (8)           | <i>nvtagpg</i> (8)           |
| <i>ping</i> (8)              | <i>pmgrd</i> (8)             | <i>poll</i> (2)              | <i>prpasswd</i> (4)          |
| <i>prpasswd</i> (8)          | <i>psradm</i> (8)            | <i>rarp</i> (1)              | <i>restore</i> (8)           |
| <i>rm</i> (1)                | <i>rmvol</i> (8)             | <i>route</i> (8)             | <i>sh</i> (1b)               |
| <i>snmp_request</i> (8)      | SSH (various)                | <i>sshd2</i> (4)             | <i>sshd2_conf</i> (4)        |
| <i>svc.conf</i> (4)          | <i>svcsetup</i> (8)          | <i>sys_attrs_alt</i> (5)     | <i>sys_attrs_bcm</i> (5)     |
| <i>sys_attrs_cam</i> (5)     | <i>sys_attrs_cfs</i> (5)     | <i>sys_attrs_clua</i> (5)    | <i>sys_attrs_clubase</i> (5) |
| <i>sys_attrs_ee</i> (5)      | <i>sys_attrs_emx</i> (5)     | <i>sys_attrs_generic</i> (5) | <i>sys_attrs_inet</i> (5)    |
| <i>sys_attrs_io</i> (5)      | <i>sys_attrs_ipv6</i> (5)    | <i>sys_attrs_lsm</i> (5)     | <i>sys_attrs_net</i> (5)     |
| <i>sys_attrs_netrain</i> (5) | <i>sys_attrs_nfs</i> (5)     | <i>sys_attrs_proc</i> (5)    | <i>sys_attrs_rdg</i> (5)     |
| <i>sys_attrs_vfs</i> (5)     | <i>sys_attrs_ufs</i> (5)     | <i>sys_attrs_vm</i> (5)      | <i>sysconf</i> (8)           |
| <i>tar</i> (1)               | <i>tcpdump</i> (8)           | <i>uucp</i> (1)              | <i>uuencode</i> (1)          |
| <i>vdump</i> (8)             | <i>vi</i> (1)                | <i>vlan</i> (7)              | <i>vlanconfig</i> (8)        |
| <i>volassist</i> (8)         | <i>voliod</i> (8)            | <i>volrestore</i> (8)        | <i>volsave</i> (8)           |
| <i>volwatch</i> (8)          | <i>vsbmpg</i> (8)            | <i>wait</i> (2)              | <i>which</i> (1)             |
| <i>wtmpfix</i> (8)           | <i>xmesh</i> (1)             | <i>ypset</i> (8)             |                              |

## Summary of Base Operating System Patches

The following sections provide brief descriptions of the changes delivered in this patch kit and in prior Version 5.1B patch kits.

The code required to make these changes may be delivered by multiple patches. As a result, you may see the same description listed under more than one patch.

Also, because Tru64 UNIX patch kits are cumulative, each new kit contains all of the fixes and enhancements that were provided in earlier kits. This can result in changes to components being made multiple times. For example, a patch may deliver several versions of the same driver. In such cases, the latest version is installed on your system.

Each patch provides fixes to subsets of the operating system. Subset names (listed in italic font in the following lists) consist of three parts; for example, for subset *OSFACCT540*, the *OSF* indicates that the subset is part of the base system, the *ACCT* indicates a category, and the *540* indicates that the subset belongs to the Version 5.1B operating system.

## New Patches

The patch summaries in this section describe defect fixes new to this release. There are no new features or enhancements added in this release.

### Patch 29002.00

#### *OSFADVFS540*

- Fixes an incomplete directory path message displayed by `vrestore (8)` while restoring device files in verbose mode.
- Fixes a problem in `vrestore (8)`, where the ASU acls are not restored for symbolic links.
- Fixes a failure in the defragment of large stripe files.

### Patch 29003.00

#### *OSFADVFSBIN540*

- Fixes the problem with `truncate` system call. When a file is being truncated in such a way that the new EOF falls in a frag, the last character at the new EOF is inadvertently overwritten with a null character.
- Fixes an issue with AdvFS performance from fragmentation of storage due to the creation of sparse files.
- Fixes a problem of kernel memory fault caused when, in certain circumstances, `rmvol (8)` and `cfsmgr (8)` are required to be running in parallel.

### Patch 29009.00

#### *OSFBASE540*

- Fixes the issue of non-root users causing cluster panic while using `sysconfig (8)`.

### Patch 29010.00

#### *OSFBIN540*

- Fixes a kernel memory fault in `ip_accessfilter`.
- Fixes a KMF problem during invalid cache memory access in the `ksm` component, which occurs when deleting and rescanning hardware devices.
- Fixes vpage array overflow issues.
- Fixes the window scale problem in TCP. The server is modified to update the window scaling option as soon as the first ACK is received from the client.
- Fixes a performance issue with `recv` system call on a socket.
- Fixes an issue in `nxm` thread creation.

### Patch 29020.00

#### *OSFCLINET540*

- Enables Tru64 xntpd to support Meinberg clocks.
- Fixes a potential security problem reported on BIND running on the HP Tru64 UNIX Operating System, where remote attackers could spoof answers returned from zones using some of the DNSKEY algorithms.
- Fixes the remotely exploitable buffer overflow problem identified in HP Tru64 UNIX NTP (Network Time Protocol) and prevents unexpected termination of an application while requesting for peer information from a remote time server.
- Fixes a potential security problem reported on BIND running on the HP Tru64 UNIX Operating System, where remote attackers could cause a DoS by sending a specific dynamic update message.
- Fixes a potential security problem that has been identified in xntpd, which may allow an attacker to cause DoS.

#### Patch 29021.00

##### *OSFCMPLRS540*

- Fixes the external timezone variable which sometimes does not get initialized.

#### Patch 29038.00

##### *OSFINCLUDE540*

- Corrects assert (3) to ensure it complies with standards.

#### Patch 29052.00

##### *OSFLSMBIN540*

- Corrects an issue with boot failure when root domain is configured under LSM. Corrects a memory leak in the LSM volobject.

#### Patch 29077.00

##### *OSFSERVICETOOLS540*

- Fixes a problem in collect (8). When running as a non-root user, collect (8) fails to gather cluster\_root domain data under AdvFS subsystem statistics.

#### Patch 29088.00.00

##### *OSFXC2SEC540*

- Fixes the incorrect behavior of dxchpwd (1X).

#### Patch 29114.00

##### *OSFSSOSSL540*

- Fixes a potential security vulnerability reported on the SSL running on the HP Tru64 UNIX Operating System, through which a remote attacker could execute arbitrary code or cause a DoS.
- Fixes a potential security vulnerability identified in the SSL running on the HP Tru64 UNIX Operating System, through which a remote attacker could cause a DoS with invalid memory access and cause the application to crash.

## Patches Delivered in Previous Kits

The following patches were first delivered in previous Version 5.1B patch kits. These patches will be installed on your system if they are newer than the last patch kit you installed.

#### Patch 28001.00

##### *OSFACCT540*

- Corrects problems with RFC1323 TCP timestamps and PAWS implementation.
- Provides an enhancement to remove drawbacks of using TCP keepalives and make LAN cluster more resilient.

## Patch 28002.00

### *OSFADVFS540*

- Corrects a time stamp problem in .
- Fixes `vrestore` to restore the file attributes properly.
- Fixes an issue with `vdump (8)`, where it gets into an infinite loop after the end of tape is reached.

## Patch 28003.00

### *OSFADVFSBIN540*

- Fixes a hang that may occur when filesystem is full and a clone fileset is created.
- Corrects an AdvFS problem that results in the `vdump` command consuming lots of CPU time while handling striped files that are less than 8 KB in size.
- Fixes a problem in AdvFS where the `bfaccess` structure is incorrectly recycled.
- Correct user visible AdvFS messages to reference "HP" instead of "Compaq".
- Provides tunable `AdvfsFragGroupDealloc` to set frag group deallocation policy.
- Fixes an issue in the AdvFS migrate code path that occurs if a file is highly fragmented and has a clone.
- Fixes a panic caused by a thread performing the migrate operation as part of a `defragment (8)` or a `migrate (8)` request.
- Fixes an issue in AdvFS resulting in a backup program receiving an incorrect sparseness map.
- Fixes an issue in the AdvFS `fs_cleanup_thread` message handling loop which, under certain circumstances, may cause it to monopolize the processor and not let other threads run on it.
- Fixes an issue with AdvFS where under certain circumstances, `vdump` routine does not return correct file extent information.
- Provides warning messages if live dumps are not collected when an AdvFS domain panics.
- Fixes an issue with AdvFS where it is not handling the truncation of sparse files properly. This results in inconsistent data.
- Fixes a problem with indexed directories that occurs when there is simultaneous access to a directory that is indexed through the AdvFS syscall interface.
- Corrects an AdvFS problem that results in the `vdump` command not backing up AdvFS files greater than 8 GB in size properly.

## Patch 28009.00

### *OSFBASE540*

- Fixes the RPC (TCP) connection reset problem when receiving XDR EOR packet without data( zero sized XDR EOR fragment).
- Fixes `find (1)` to match the pathname, if the base file name of the pathname matches the pattern specified in `-name`, even when the pathname has trailing slash(es).
- Fixes a problem in `scanf` family of functions when format specifier is `%%`.
- Fixes an issue with `fopen()`, `fdopen()` and `popen()`.
- Fixes a failure in `printf` family of functions.
- Fixes a problem in `/sbin/init.d/cron` which prevented the `cron` daemon from restarting in Japanese locale.

- Provides additional Fibre Channel HBA port type definitions in the FC-HBA header file, `hbaapi.h`.
- Fixes a problem in `printf` family of functions when number grouping is requested.
- Corrects DST starting 3 December 2006 for Western Australia per the Public Domain time zone source kit (`tzdata2007a`).
- Fixes a file descriptor leak in `pthread` library.
- Update `newfs` to not set `disklabel` on `newfs` failure.
- Fixes problems with the IBM Tivoli Storage Manager and the `libpac1.so` library routines.
- Fixes a problem in `return` built-in command of `ksh`, when it is used in a sub-shell.
- Addresses the issue of `sizer v` not displaying the updated HP re-brand information when patch kit is installed on alternative root with `-root` option of `dupatch`.
- Removes duplicate lines from `df` output.
- A potential security vulnerability has been identified in the `ps` and `w` commands on the HP Tru64 UNIX Operating System. The `ps` and `w` commands can be used to disclose environmental variable and argument information that might be exploited by a local, authorized user.
- Corrects a problem in `lpd` resulting in creation of many `lpd` child processes when the size of `/etc/printcap` exceeds 8192 bytes.
- Corrects New Zealand DST changes starting September 2007.
- Fixes an issue in `/sbin/mountroot` script which, under DMS (Dataless Management Service), would attempt a mount update operation on the NFS root filesystem.
- Eliminates `auditd` zombies resulting from combination of `syslog` failure and `auditd` overflow handler.
- Fixes a problem in `pthread_mutexattr_getprotocol()`.
- Fixes `dumpfs` to report the proper error message. Updates `dumpfs` to handle filesystem images.
- Corrects South East Australia DST changes starting April 2008.
- Fixes an issue with `strip` command.
- Fixes incorrect `exit/continue` logic in `mailsetup`.
- Fixes a problem in `/sbin/init.d/envmon` which prevented the `envmon` daemon from starting in Japanese locale.
- Fixes a problem in `/sbin/init.d/smauth` and `/sbin/init.d/smsd` which prevented the `smauth` and `smsd` daemons from starting or restarting in Japanese locale.
- Mitigates a queue control file version incompatibility that can lead to mail messages being orphaned
- Updates the Canada, Bahamas, Bermuda, Brazil, and Uruguay DST changes 2007 per the latest PD time zone source kit (`tzdata2007a`).
- Fixes a problem in which the `disklabel` utility sets the partition size incorrectly on a CDFS (ISO 9660) ISO image file when the size is larger than 4 GB.
- Fixes the DNS problem, where DNS client loops with unknown incoming UDP packets.
- Adds support in the `audit_tool` utility, in "brief" mode, to map a socket.
- Adds support in the `audit_tool` utility for new `kmodcall` opcodes.
- Changes the string from "Compaq Computer Corporation" to "Hewlett-Packard Company" in `sendmail` file `makeinfo-dec.sh`.
- Fixes the issue of displaying HP brand information when patchkit is installed with `-root` option of `dupatch`.
- Fixes a problem in handling invalid inputs in the `libc inet_network` routine.
- Modifies `ksh (1)` built-in test to evaluate string expressions as per POSIX Standard.
- Provides Venezuela & Argentina with new DST changes starting December 2007.

- Removes the code for `smmosp` user and group creation.
- Fixes an issue in `strncmp`.

## Patch 28010.00

### *OSFBIN540*

- Fixes a potential system panic from an illegal memory reference due to a race condition in the class scheduling code.
- Fixes consumption of processor resources for an extended period creation of a core file for processes with large address spaces, thereby avoiding blocking other threads of execution and leading to a time out panic.
- Fixes a deadlock due to respective buffers sharing a physical page concurrent AIO operations to a shared memory region from multiple processes.
- Fixes a system panic when `tcbhashsize` tunable is modified on a running system with lock mode 4 enabled.
- Fixes for Mobile IPv6 TAHi conformance test failures.
- Fixes a Path MTU problem in network stack.
- Fixes double unlocking of `netisr_lock` in `netisr` thread.
- Resolves synchronization issues surrounding concurrent references to shared memory.
- Fixes a SACK (Selective Acknowledgement) validation issue with TCP.
- Resolves a race condition between process exit and `coredump`.
- Fixes a problem in the FIFO code where usage of `fuser(2)` subsequent to opening a FIFO may result in a kernel panic with the panic string `VREF: Invalid v_usecount transition`.
- Fixes a problem wherein the signal mask was not getting restored after returning from `sigsuspend`.
- Fixes potential hang or panic resulting from blocking under `select_enqueue()`.
- Turns on the TCP Selective Acknowledgment (SACK) option by default.
- Corrects a problem with FIFO code where `stat(2)` or `fstat(2)` operation on a FIFO returns `NODEV` in the `st_dev` field of the `stat` structure.
- Corrects an issue that can lead to a kernel memory fault panic when locking file backed memory on a NUMA system.
- Corrects problems with RFC1323 TCP timestamps and PAWS implementation.
- Fixes a timing window in the VM subsystem and prevents a race condition between the page fault and the pageout paths.
- Fixes an issue with NFS `vnode` reclaim path in CFS (Cluster File System) environment.
- Fixes loss of multicast packet reception after `netrain` failover.
- Allows filtering of `dup2` syscalls while audit object selection is enabled.
- Provides an enhancement to remove drawbacks of using TCP keepalives and make LAN cluster more resilient.
- Corrects a potential security vulnerability whereby, under certain circumstances, a local authenticated user could cause a Denial of Service.
- Fixes TCP SACK unaligned access panic.
- Improves performance of `ufs extendfs` operation on a mounted filesystem.
- Introduces a new `sysconfig` tunable to control the behavior of the `O_APPEND` flag with respect to the `pwrite(2)` system call.
- Fixes a problem with the kernel traffic monitoring thread.
- Fixes a problem in FFM unmount code where, under certain circumstances two or more unmount threads race and one thread can return without releasing the NFS `lock_for_exported` results in a hang during unmount or shutdown.



- Corrects `fixfdmn` to display proper error messages when `open()` or `write()` calls to the device fail.
- Fixes kernel memory fault panic in network stack.
- Corrects an issue with the TCP congestion control algorithm in Tru64 UNIX.
- Corrects a problem with the Tru64 NFS server where, under certain conditions, directories may not be visible to clients.
- Fixes a problem with `read` of `/dev/mem` which resulted in `netstat` hang.
- Add support in cluster alias to handle socket unlisten.
- Fixes an improper reference in `anon_dup` routine.
- Fixes an issue with handling page-faults inside kernel.
- Fixes an issue in FIFO open, where under certain circumstances, applications using FIFOs may hang.
- Fixes an issue with `table()` system call on a NUMA based system with a memoryless RAD configuration.
- Provides event markers to track missing `binary.errlog` events.
- Fixes a lock wait panic in the scheduling subsystem.
- Fixes a synchronization issue between `fuser()` and `exec_close_files()`.
- Fixes problem in reporting tcp sacks.
- Revises the UBC page allocation algorithm to consider wired pages.
- Fixes for IPv6 TAHI conformance test failures.
- Resolves a synchronization issue in anon subsystem.
- Corrects an improper data access issue in the audit subsystem when auditing system calls that operate on UNIX domain sockets.
- Fixes a locking issue in the audit subsystem.

## Patch 28020.00

### *OSFCLINET540*

- Fixes the RPC (TCP) connection reset problem when receiving XDR EO packet without data(i.e zero sized XDR EOR fragment).
- Fixes the problem with `rcp` where it fails if the remote username contains a dollar(\$) sign.
- BIND 9.2.8 release.
- Provides envmon subsystem enhancements on DS-25 and ES-45 plat
- Corrects the security vulnerability, SSRT 080058: DNS Cache Poisoning.
- BIND 9.2.8-P1 release.
- Fixes a problem with `inet` script looping.

## Patch 28021.00

### *OSFCMPLRS540*

- Fixes RPC (TCP) connection reset problem when receiving XDR EOR packet without data(zero sized XDR EOR fragment).
- Fixes a problem in `scanf` family of functions when format specifier is `%%`
- Fixes an issue with `fopen()`, `fdopen()`, and `popen()`
- Fixes a failure in `printf` family of functions.
- Fixes a problem in `printf` family of functions when number grouping is requested.
- A potential security vulnerability has been identified in the `ps` and `w` commands on the HP Tru64 UNIX Operating System. The `ps` and `w` commands can be used to disclose environmental variable and argument information that might be exploited by a local, authorized user.
- Fixes the DNS problem, where DNS client loops with unknown incoming UDP packets.

- Fixes a problem in handling invalid inputs in the `libc inet_network` routine.
- Fixes an issue in `strncmp`.

### Patch 28028.00

#### *OSFENVMON540*

- Fixes a problem in `/sbin/init.d/envmon` which prevented the `envmon` daemon from starting in Japanese locale.
- Fixes a problem in `/sbin/init.d/smauth` and `/sbin/init.d/smsd` which prevented the `smauth` and `smsd` daemons from starting or restarting in Japanese locale.

### Patch 28034.00

#### *OSFHWWBASE540*

- Corrects firmware problem with HSZ70.
- Fixes an error handling problem with list traversal in `netstat(1)`.

### Patch 28035.00

#### *OSFHWWBIN540*

- Fixes an issue with tape devices where a rewind may occur when using a no rewind device special file. This issue is more likely to occur if the `sysconfig` variable `cam_ccfg_aa_enable` in the `cam` module is set to 1, the default setting, and the device reports that it supports asymmetric logical unit access in the inquiry data.
- Fixes a problem with the CAM configuration driver by releasing the memory allocated after completing the `REPORT TARGET PORT GROUP` command.
- Adds a new feature to the CAM disk driver to dynamically recognize and use new paths while the device is in use.
- Fixes potential hang in tape backup applications.
- Resolves a locking issue within the USB subsystem that can lead to a system failure.
- Fixes incorrect CAM status from `aha_chim` driver during bus reset processing.
- Adds support for 2TB LUNs.
- Fixes an issue with handling `ssm` and `l3gh`.
- Prevents simple lock timeout panic in `emx` driver during mailbox commands.
- Provides `envmon` subsystem enhancements on DS-25 and ES-45 platforms.
- Fixes a panic in `alt` driver.
- Corrects panic seen during hardware registration.
- Corrects a rare circumstance where an IO can erroneously be returned as successful when it has failed.
- Provides event markers to track missing `binary.errlog` events.
- Corrects a Kernel Memory Fault panic in the `emx` driver.
- Enables capturing of additional state during a system failure to facilitate the analysis of virtual memory related issues.
- Fixes a Kernel Memory Fault panic in the `KZPEA` driver.
- Fixes a timing issue during boot process.
- Fixes an issue that caused kernel to read thermal data from power supplies across partition boundaries.
- Fixes Kernel Memory Fault on inquiry to a non existing LUN connected through `KZPEA`.
- Fixes "Freed CCB" panic and lock issues.
- Fixes a panic in `bcm` driver.
- Provides 2TB LUN support in disk recovery.
- Fixes a KMF during boot when `KMEM_DEBUG_PROTECT` is enabled.

- Fixes a problem with read of /dev/mem which resulted in netstat hang.
- Prevents monitoring thread issuing IO when disk is in recovery.
- Provides event markers to track missing binary.errlog events.
- Fixes mistaken report of user id change in audit data for mach trap.
- Fixes an issue with usage of GH regions and chunks.
- Fixes an issue in pointer synchronization for USB hub devices.
- Fixes potential hang in tape backup applications.
- Fixes a problem in getsysinfo.
- Fixes an issue with the handling of Segmented Shared Memory (SSM) mapped for read-only access.

#### Patch 28038.00

##### *OSFINCLUDE540*

- Fixes a problem with L\_tmpnam macro.

#### Patch 28039.00

##### *OSFINET540*

- BIND 9.2.8 release.
- Fixes a RFC1323 TCP timestamp display problem with tcpdump.
- A potential security vulnerability has been reported on the HP Tru64 operating system running BIND. The vulnerability is remotely exploitable and may result in DNS cache poisoning.
- Fixes various functionality issues in Sysman mail application.
- BIND 9.2.8-P1 release.

#### Patch 28045.00

##### *OSFKTOOLS540*

- Provides a command-line option (-ignore\_pset) to lockinfo command to print the lock details of all processors (-percpu) ignoring processor set boundaries.
- Fixes a problem wherein patch installation fails to remove temporary files.
- Corrects a problem with the kdbx debugger that caused the tool to display incomplete socket/pcb addresses.
- Fixes the kdbx extensions ilock, slock, export, netstat, abscallout, rpb, ofile, and plist by synchronizing with the current release.

#### Patch 28046.00

##### *OSFLAT540*

- Fixes an unaligned memory access in LATCP.

#### Patch 28050.00

##### *OSFLIBA540*

- Update newfs to not set disklabel on newfs failure
- Upgrades libots3.a to V2.1-121

#### Patch 28051.00

##### *OSFLSMBASE540*

- Fixes a problem with LSM that renders it incapable of recognizing disk clones correctly
- Fixes a problem in vold that can cause voldisksetup to hang while initializing KZPCC disks.

- Updates `volstat` utility and kernel to report cluster-wide LSM statistics.
- Fixes a diskgroup deport problem during multiple plex detaches.

## Patch 28052.00

### OSFLSMBIN540

- Fixes a `simple_lock_fault` in LSM's dirty region log code which is seen under a low memory condition.
- Fixes the `clsm` checksum error seen with a disappearing `vold` while `volrestore` is run.
- Adds support into LSM for 2TB LUNs.
- Updates `volstat` utility and kernel to report cluster-wide LSM statistics.

## Patch 28054.00

### OSFMANOP540

- Updates the following reference pages:

|                                  |                                      |                                   |                                    |
|----------------------------------|--------------------------------------|-----------------------------------|------------------------------------|
| <a href="#">aliases(4)</a>       | <a href="#">advfsd(8)</a>            | <a href="#">disklabel(8)</a>      | <a href="#">dsfmgr(8)</a>          |
| <a href="#">dumpfs(8)</a>        | <a href="#">fcntl(2)</a>             | <a href="#">fread(3)</a>          | <a href="#">freezefs(2)</a>        |
| <a href="#">gethostbyaddr(3)</a> | <a href="#">gethostbyname(3)</a>     | <a href="#">getlogin(2)</a>       | <a href="#">getrlimit(2)</a>       |
| <a href="#">iconv_intro(5)</a>   | <a href="#">ksh(1)</a>               | <a href="#">lockinfo(8)</a>       | <a href="#">mailstats(8)</a>       |
| <a href="#">mail.local(8)</a>    | <a href="#">mailq(1)</a>             | <a href="#">makemap(8)</a>        | <a href="#">newaliases(1)</a>      |
| <a href="#">netstat(1)</a>       | <a href="#">praliases(1)</a>         | <a href="#">proplist(4)</a>       | <a href="#">ps(1)</a>              |
| <a href="#">sendmail(8)</a>      | <a href="#">sftp2(1)</a>             | <a href="#">sialog(4)</a>         | <a href="#">signal(2)</a>          |
| <a href="#">ssh2(1)</a>          | <a href="#">ssh-agent2(1)</a>        | <a href="#">ssh2_config(4)</a>    | <a href="#">sys_attrs_advfs(5)</a> |
| <a href="#">sys_attrs_cam(5)</a> | <a href="#">sys_attrs_generic(5)</a> | <a href="#">sys_attrs_inet(5)</a> | <a href="#">sys_attrs_vfs(5)</a>   |
| <a href="#">sys_attrs_vm(5)</a>  | <a href="#">useradd(8)</a>           | <a href="#">vacation(1)</a>       | <a href="#">vfast(8)</a>           |
| <a href="#">vipw(8)</a>          | <a href="#">w(1)</a>                 |                                   |                                    |

## Patch 28055.00

### OSFMANOS540

- Updates the [find\(1\)](#) reference page.
- Updates the [mlock\(3\)](#) reference page.
- Updates the [sys\\_attrs\\_generic\(5\)](#) and [sys\\_attrs\\_advfs\(5\)](#) reference page.
- Corrects the value of `vm_ubcdirtypercent` to 40 from 10 in the reference page of [sys\\_attrs\\_vm\(5\)](#).
- Updates the [sys\\_attrs\\_inet\(5\)](#) reference page.
- Updates [mount\(2\)](#), [unmount\(2\)](#), [mount\(8\)](#), and [unmount\(8\)](#) reference pages.
- Updates the following reference pages:

|                                  |                                  |                              |                               |
|----------------------------------|----------------------------------|------------------------------|-------------------------------|
| <a href="#">aliases(4)</a>       | <a href="#">advfsd(8)</a>        | <a href="#">disklabel(8)</a> | <a href="#">dsfmgr(8)</a>     |
| <a href="#">dumpfs(8)</a>        | <a href="#">fcntl(2)</a>         | <a href="#">fread(3)</a>     | <a href="#">freezefs(2)</a>   |
| <a href="#">gethostbyaddr(3)</a> | <a href="#">gethostbyname(3)</a> | <a href="#">getlogin(2)</a>  | <a href="#">getrlimit(2)</a>  |
| <a href="#">iconv_intro(5)</a>   | <a href="#">ksh(1)</a>           | <a href="#">lockinfo(8)</a>  | <a href="#">mailstats(8)</a>  |
| <a href="#">mail.local(8)</a>    | <a href="#">mailq(1)</a>         | <a href="#">makemap(8)</a>   | <a href="#">newaliases(1)</a> |

|                          |                              |                           |                            |
|--------------------------|------------------------------|---------------------------|----------------------------|
| <i>netstat</i> (1)       | <i>praliases</i> (1)         | <i>proplist</i> (4)       | <i>ps</i> (1)              |
| <i>sendmail</i> (8)      | <i>sftp2</i> (1)             | <i>sialog</i> (4)         | <i>signal</i> (2)          |
| <i>ssh2</i> (1)          | <i>ssh-agent2</i> (1)        | <i>sshd2_config</i> (4)   | <i>sys_attrs_advfs</i> (5) |
| <i>sys_attrs_cam</i> (5) | <i>sys_attrs_generic</i> (5) | <i>sys_attrs_inet</i> (5) | <i>sys_attrs_vfs</i> (5)   |
| <i>sys_attrs_vm</i> (5)  | <i>useradd</i> (8)           | <i>vacation</i> (1)       | <i>vfast</i> (8)           |
| <i>vipw</i> (8)          | <i>w</i> (1)                 |                           |                            |

## Patch 28062.00

### *OSFNFS540*

- Fixes an issue with NFS `rpc.lockd` in the lock reclaim path.
- Fixes the `mountd` problem when user supplies wrong port number.
- Fixes an issue with using LDAP backed netgroups with NFS `mountd`.
- Fixes a problem where `mountd` exits due to segmentation violation.
- Fixes memory leak issues with NFS `rpc.lockd` daemon.
- Fixes a problem with `lockd` where, under certain circumstances it may terminate with a core dump because of incorrect usage of memory on which `malloc()` has been performed.

## Patch 28069.00

### *OSFPRINT540*

- Fixes an issue with `lpr(1)`, when the `/etc/printcap` file has syntax error.
- Added support for new printers.
- Corrects a problem in `lpd` resulting in creation of many `lpd` child processes when the size of `/etc/printcap` exceeds 8192 bytes.

## Patch 28077.00

### *OSFSERVICETOOLS540*

- Corrects `collect(8)` to gather `cluster_root` domain information, when run as non-root user.
- Fixes an issue in `collect` for data collection problem for case where an AdvFS domain information is not accessible.
- Solves a problem in `collect` utility, where `collect` exits with the error message "popen() failed: Too many open files", when run in historical mode.

## Patch 28079.00

### *OSFSYSMAN540*

- Corrects a potential security vulnerability that may lead to improper file access.
- Fixes a problem in `/sbin/init.d/envmon` which prevented the `envmon` daemon from starting in Japanese locale.
- Fixes a problem in `/sbin/init.d/smauth` and `/sbin/init.d/smsd` which prevented the `smauth` and `smsd` daemons from starting or restarting in Japanese locale.
- Corrects a potential security vulnerability in `dop` where, under certain circumstances, a user could potentially execute privileged code.

## PATCH 28085.00

### *OSFX11540*

- This fix addresses a X-Motif list widget problem while selecting an item from a long list.

## PATCH 28102.00

### *OSFCDSABASE540*

- Fixes a problem with CDSA configuration where `mod_install` program can core dump.

## PATCH 28107.00

### *OSFLDPAUTH540*

- Fixes ldapcd daemon to service the requests faster after failover from primary Active Directory to Backup Active Directory.
- Fixes the locale problem in ldapcd daemon script.

## PATCH 28113.00

### *OSFSSHBASE540*

- Corrects a potential security vulnerability that has been identified with SSH running on the HP Tru64 UNIX Operating System. Vulnerability could be exploited to allow remote unauthorized access to sensitive information.
- Updates the SSH client to use protocol version 2.
- Fixes wildcard matching and globing in `scp2/sftp2 ls -l` command.
- A potential security vulnerability has been identified in the SFTP server (`sftp-server`) component of SSH 3.2.0 and earlier running on HP Tru64 UNIX versions 5.1B-3 and 5.1B-4. The vulnerability could be exploited remotely to allow an authorized remote attacker to execute arbitrary code or cause a denial of service (DoS).

## Patch 27001.00

### *OSFACCT540*

- Corrects the action of the `dodisk` command to skip the commented file systems contained in the `/etc/fstab` file.
- Fixes `acctcom` to exit with proper error message when used with invalid user ID and group ID.
- Causes the file protections of `/var/adm/pacct` and ownership of `/var/adm/wtmp` to act as expected by the accounting utilities.
- Fixes the `fwtmp` command so it does not display invalid (negative) PIDs when the number of decimal digits of PID value exceeds 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.
- Corrects an error in the script in `lastlogin.sh`.
- Makes start up scripts in `/sbin/init.d` world readable.
- Corrects the following problems found in accounting commands:
  - Resolves the differences in the CPU time and connect time found during the conversion of accounting reports from ASCII format to binary and again back to ASCII.
  - Resolves the differences in CPU time found in the output of the `acctcom` and `acctmerg` commands for the same input file.
  - Fixes the way accounting files are referenced using CDSLs.
  - Corrects the display of the header from `acctcom` when accounting is first started.
  - Corrects an error message during execution of the `runacct` command.

- Enables acctcom to read more than one input file.
- Enables acctcom to work with pipes.
- Fixes the acctcom command with respect to the display of tty lines.
- Corrects a problem in which accounting reports show the wrong connection time for the users who remain logged in during the execution of runacct.
- Enables the acctcon1 command to calculate the connect time for local logins in case of run level changes.

## Patch 27002.00

### *OSFADVFS540*

- Corrects the following problems with the fixfdmn utility:
  - Fixes fixfdmn so that it will continue when it finds more than one root tag file, rather than exiting with the message “Unable to continue, more than one root tag file.”
  - Fixes a potential memory fault while running the fixfdmn.
  - Fixes a problem in which fixfdmn does not properly handle domains with multi-page RBMT files.
- Fixes the incorrect number of tag pages being returned by some AdvFS user commands.
- Fixes vrestore command to restore file and directory attributes by default and addresses security issues during the restoration of those attributes.
- Removes memory leaks from the AdvFS salvage utility.
- Fixes rare problem in which the vdump/vrestore commands hang when requesting the next tape.
- Enhances the AdvFS rmvol utility to allow multiple volumes to be removed with one command.
- Fixes a problem that causes the defragment command to fail with the error message “defragment: Can't allocate memory”
- Corrects return values from the vfast utility.
- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file access.
- Provides enhancements for file system suitlets.
- Improves AdvFS informational messages as follows:
  - Advscan reports if a domain has all of its volumes, but they are stored in a different directories. This scenario will cause mount to fail.
  - The AdvFS I/O error message includes the location of a file that will help users to translate the error number into an error message.
- Prevents erroneous display of DMAPI messages while using the showfile command.
- Fixes a problem in which the verify utility core dumps if it encounters a specific type of metadata inconsistency.
- Improves the information provided in an error message that is displayed when the chfile command to display an informative error message if it fails while trying to enable data logging.
- Modifies the nvbmtpg utility to display all the data in a specified mcell only if the -v option is selected.
- Fixes a problem where defragmentation can fail if the process can not obtain enough memory.
- Provides support for Smart Array disk controllers. Without this patch, if the Smart Array product is installed on the system, the SysMan Station hardware view will fail to operate.
- Corrects the error message returned when trying to migrate striped files when the -s option is omitted.

- Makes the following changes to the `vdump`, `rvdump`, and `vrestore` commands:
  - Causes `vdump` and `rvdump` to report when all hard link siblings cannot be archived through the specified path and causes them to correct the bytes to backup estimate calculation when hard links are archived.
  - Causes `vdump` and `vrestore` to act as expected upon receiving an interrupt (^C).
  - Fixes `vdump` and `vrestore` to pick up correct messages in all locales.
  - Causes `vdump` to avoid some unnecessary function calls, thereby allowing faster `vdumps`.
  - Fixes `vrestore` to display bit file attributes with the `-l` option.
  - Prevents `vrestore` from failing during a remote system call.
  - Causes `vrestore` to display a file and directory name along with the error message when the command fails to set a property list.
  - Prevents `vrestore` from dumping core when a tape has a smaller blocksize than expected.
  - Allows `vrestore` to handle no-rewind tapes properly.
  - Lets `vrestore` read environment variables for a user-defined device name.
  - Allows attributes to set to the top level-directory.
  - Disallow block sizes greater than 64 K. (This was a new feature added in Patch Kit 2 which caused forward compatibility problems.)
  - Causes files of less than 512 bytes to be properly backed up.
  - Corrects a cluster `drd` bandwidth problem that occurs when a restore is attempted via a private tape drive that is not on the current node.
  - Corrects a condition that could cause the misinterpretation of the archive version in certain locales, causing `vrestore` to abort.
- Corrects several problems with the `fixfdmn` utility:
  - Shortens the time it can take to repair large AdvFS file domains.
  - Improves the capability of `fixfdmn` to fix invalid extent data.
  - Improves the secure handling of temporary files by `fixfdmn`.
  - Allows `fixfdmn` to repair a rare on-disk structure problem with directories. This condition only shows up on a domains with multiple volume in which property lists are used.
  - Improves the ability of `fixfdmn` to fix on-disk structure version 3 domains.
  - Fixes several rare cases where `fixfdmn` could either fail to correct a domain or incorrectly make changes to a valid domain.
- Fixes many small problems with the `dsfmgr` command.
- Provides protection against a class of potential security vulnerabilities called buffer overflows. Buffer overflows are sometimes exploited in an attempt to subvert the function of a privileged program and possibly execute commands at the elevated privileges if the program file has the `setuid` privilege. This patch allows a system administrator to enable memory management protections that limit potential buffer overflow vulnerabilities.
- Corrects a problem with the `AdvFSstat` command printing negative values for statistics that are over 10 decimal digits long.
- Corrects several rare cases in which `fixfdmn` either did not fix, or incorrectly fixed data inconsistencies on AdvFS file domains. On large AdvFS file domains, this patch shortens the time it takes to repair them.
- Causes `vdump` and `rvdump` to report when all hard link siblings cannot be archived through the specified path.
- Causes `vdump` and `rvdump` to correct the bytes-to-backup-estimate calculation when hard links are archived.
- Enhances `/sbin/AdvFS/tag2name` to print out the name of the associated directory, given the tag of an index file.



- Allows fixfdmn to repair a rare on-disk structure problem with directories. This condition only shows up on a domains with multiple volume in which property lists are used.
- Corrects several very rare cases where fixfdmn incorrectly fixes certain AdvFS file domains.
- Causes the vdump -C option to display compression ratio.
- Causes vrestore to display error message for end of tape.
- Identifies non-tape target device during rvdump and vrestore actions.
- Improves the ability of fixfdmn to fix on-disk structure version 3 domains.
- Corrects the value nvbmtpg prints for fs\_stat.st\_unused\_1.
- Ensures that a stripe does not succeed if a write precedes it.
- Corrects the RBMT free mcell count for multipage RBMTs as displayed by nvbmtpg.
- Allows the maximum directory entry name size in the vods tools.
- Corrects several rare cases where fixfdmn incorrectly fixes certain AdvFS file domains.
- Modifies defragcron.sh to avoid accidentally deleting the /dev/null.
- Enables the vfast utility to run on a standalone system.
- Fixes the vfast utility error "vfast: cannot get frag list; 14 - Bad address"
- Fixes problems with vdump to resolve filesets mounted on a CDSL.
- Fixes problems with vrestore to restore appropriate ACLs on the directories and sub-directories.
- Corrects a problem that occurs when fixfdmn is run with an invalid fileset argument and the user is prompted if exit or run it against all the filesets in the domain. If exit fixfdmn is selected, changes may be made to the domain, which could cause problems.
- Modifies vdump and vrestore to support autoloader or manual loader when the dump size exceeds single tape size
- Modifies vdump and vrestore so they can be run as cron or batch jobs.
- Fixes a problem in which the SysMan menu Manage AdvFS File Systems application does not retrieve the volume attributes correctly.
- Enforces the validation of incorrect date which would go unchecked with the salvage -d option
- Causes vdump to display the correct number and percentage of bytes to back up when a directory with symbolic links is dumped.
- Modifies vfast so it displays the correct statistics data for the -L extents option.
- Causes vrestore to display the correct number of volumes on which the file is striped and to display a new message when the striped file is restored on a single volume.
- Fixes a situation in which vdump exits with message "Not enough memory to generate link table."
- Fixes problems with vrestore that wrongly sets ownership/permissions on the top level directory.
- Fixes a segmentation fault by the salvage utility.
- Fixes a race condition between two threads in defragment to prevent it from looping indefinitely.

## Patch 27003.00

### *OSFAdvFSBIN540*

- Fixes a problem that causes the first command that accesses an AdvFS domain to fail if the domain had a full RBMT when mounted.
- Fixes a deadlock that can happen during failover of global root and var file systems when vfast is enabled on them.
- Fixes a single thread hang.
- Fixes a potential deadlock hang between a migration and a flush on a file.

- Corrects AdvFS quota handling and enforcement to prevent EIO and false out-of-space conditions and account for directory index blocks.
- Corrects the following problem with read-only and dual mounted file systems: If the system crashes after a read-only dual mount, the next time the system is booted the AdvFS domain panics on the file systems that were previously mounted as read-only dual mount.
- Fixes a kernel memory fault from `imm_page_to_sub_xtnt_map()`.
- Fixes AdvFS AIO write error paths so the I/O completion steps are not repeated.
- Fixes a hang that occurs in file systems between racing memory-mapping threads.
- Turns a potential AdvFS panic into a domain panic.
- Corrects a problem with the AdvFS bitfile state where, under certain circumstances, the state of the bitfile is changed without holding the lock that protects the state.
- Fixes an issue with `vfast` in which a `vfast` thread might cause a "kernel memory fault" panic because of a race condition.
- Fixes a potential panic on an active domain doing heavy I/O while trying to create a new fileset.
- Fixes problems that occurs when volume expansion (`mount -u -o extend`) races with other code.
- Fixes a rare simple lock timeout during domain deactivation.
- Fixes a potential hang in AdvFS.
- Converts non-severe system panic calls to AdvFS domain panic calls while renaming files under AdvFS.
- Corrects a problem in which an ENOSPC error is returned on a fragmented AdvFS file system even though space is available. On a cluster, this can lead to a CFS WRITE ERROR.
- Fixes a premature out-of-space condition that can occur as a result of repetitively extending the size of the volume.
- Fixes resource leaks seen after a device file gets revoked.
- Fixes a delete failure due to a "Disc quota exceeded (EDQUOT)" condition.
- Fixes a data inconsistency that could result from the failure of an I/O request at the CFS Server due to exceeding the fileset quota.
- Improves AdvFS informational messages as follows:
  - 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 users to translate the error number into an error message.
- Prevents race conditions that could cause a kernel memory fault while doing a migrate and a `rmvol` on a striped file.
- Modifies the CFS flushing behavior during an `rmvol`.
- Helps reduce the size of extent maps of clone files cases where the original is modified extensively under direct I/O.
- Enhances `/sbin/AdvFS/tag2name` command to print the name of the associated directory when given the tag of an index file.
- Corrects a potential problem with modifying files via direct I/O when there is a clone fileset.
- Fixes a race during AdvFS volume removal that can cause a panic in the `bs_osf_complete()` routine.
- Fixes a problem when monitoring I/O via `AdvFSstat`.
- Changes the behavior of `migrate_normal` and `migrate_stripe` when migrating an original file that has a clone. If the clone was marked out of sync, migrate could come back with `E_CLONE_OUT_OF_SYNC` even though the migrate succeeded. Now this case is caught, and handled.

- Fixes a problem in which a domain panic may occur in `idx_lookup_node_int` or `bs_frag_dealloc` under heavy file system activity, generating one of the following messages:  
`idx_lookup_node_int: bs_refpg failed`  
`bs_frag_dealloc: rbf_pinpg (4) failed, return code = -1035`
- Fixes a part of AdvFS migration code in order to prevent `rmvol "Can't remove volume"` error.
- Replaces the system panics caused by `"Can't clear bit twice"` with a domain panic.
- 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 a hang may occur when a `rmvol` operation is performed after a cluster node failure during `volmigrate`, `volunmigrate`, or `frag` file migration.
- Fixes a problem in which a crash occurs when an AdvFS file system reports I/O errors and enters into a domain panic state. The AdvFS error cleanup would panic on an invalid pointer and report an `"invalid memory read access from kernel mode"` panic message.
- Fixes a standard violation on AdvFS.
- Helps prevent kernel memory faults in AdvFS.
- Fixes a problem where `gh_min_seg_size` could not be set below 8M.
- Corrects a race condition in AdvFS in which it avoids a potential stranded log record in memory that does not get out to disk.
- Prevents a potential hang during a reboot after a recent domain panic.
- Prevents a panic in a cluster that when a fileset mounted `-o dual` is failed over or unmounted during a shutdown.
- Fixes a possible kernel system hang in `vfast` when shutting down or rebooting the system.
- 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.
- Improves the performance of AdvFS under heavy I/O loads.
- Avoids 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`.
- Adds defensive programming to `stat.h` to avoid `stat.h` getting confused if one of its internal temporary `#defines` is defined before `stat.h` is processed.
- Corrects an internal AdvFS check that was always returning true.
- Improves the scalability and performance of AdvFS.
- Allows AdvFS to record if a domain panic has occurred, even if a system panic results.
- Replaces two potential panics in AdvFS with domain panics.
- Provides scalability improvements to AdvFS that will help reduce lock contention and improve performance.
- Fixes an AdvFS path that can cause a panic in the `advfs_page_busy()` routine.
- Fixes a deadlocking problem in the kernel where a file open on a clone could hang when ACLs are enabled.
- Fixes a hang that can occur during the renaming of an AdvFS file.
- Displays the correct error message for `freezefs -q` on a non-AdvFS file system.
- Adds comment to reserve `0x10000000` and `0x20000000` for AutoFS flags.
- Prevents a `O_DSYNC` write failure under the following situation:
  1. The user creates a new file.
  2. Closes the file.
  3. The `vnode` for the file is recycled.
  4. Reopens the file with the `O_DSYNC` flag.
  5. Writes to the file, overwriting already allocated storage.

6. The write from step 5 returns to the application.
  7. The system crashes.
- Corrects the NFS server's handling of files open for direct I/O.
  - Fixes an internal problem in the kernel's AdvFS, UFS, and NFS file systems where extended attributes with extremely long names (greater than 247 characters) could not be set on files. The new limit is 254 + a null string terminator.
  - Corrects a problem that can produce the following symptoms:
    - When producing a clone of a file (that is, a fileset has been cloned and one of its files is being written to) or when migrating a file (defragment, migrate, balance, rmvol), directories lose attributes and are seen as files, and files lose their ACLs and other attributes under the following circumstances:
      - On multi-volume domains with a volume that is out of space
      - By the use of a property list on a multi-volume domain
    - The rmvol command enters into an infinite loop while trying to move a file from one volume to another.
  - Fixes a problem where threads can hang when performing a malloc() function.
  - Prevents a kernel memory fault panic in \_OtsMove when going through the fs\_read() routine.
  - Prevents a potential hang during a umount if a domain\_panic has been encountered.
  - Provides a workaround for a domain panic when a data inconsistency in the deferred-delete list of an AdvFS file system is detected.
  - Corrects idle-wait time accounting within the AdvFS file system, as reported by vmstat -w.
  - 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).
  - Removes the obsolete function bs\_bfdmn\_flush\_all.
  - Prevents a potential unaligned memory crash when ACLs are on.
  - Eliminates the lock\_terminate: lock owned panic.
  - Fixes a problem where data from an AdvFS file with a frag could be written to an incorrect location if an NFS client grew the file.
  - Fixes an AdvFS asynchronous direct I/O problem that can cause a thread to hang.
  - Fixes a problem encountered where a truncated AdvFS file erroneously zeros data for the remaining leading segment of the file.
  - Corrects a condition that causes a panic resulting from a kernel memory fault in access\_invalidate.
  - Corrects a problem in mount or domain activation after a panic, where a fileset (domain) cannot be mounted without running fixfdmn.
  - Improves performance for CFS filesets mounted with the server\_only option. A log sync for create transactions is not needed for such filesets.
  - Fixes a cluster panic with the following error message:
 

```
panic (cpu 3): ics_unable_to_make_progress: heartbeat checking blocked
```
  - Fixes an rmvol E\_PAGE\_NOT\_MAPPED error.
  - Eliminates an ENO\_MORE\_BLKs error seen performing a copy-on-write procedure to a clone file while an rmvol operation is in progress.
  - Increases from 4 to 6 the number of pages that can be pinned at deletion time.
  - Improves the informational messages returned by a few domain panic strings.
  - Fixes an error in some sections of code that get E\_PAGE\_NOT\_MAPPED errors when the code expected the page to exist.
  - Closes a small race accessing internal data structure in AdvFS.
  - Prevent a potential panic when AdvFS looks up a file name.
  - Corrects the following problems in AdvFS write logic:

- A mismatch between the value reported by the write system call and the number of bytes written.
- Unavailable and unused storage.
- Causes the deallocation of preallocated storage that the caller is not using.
- Improves the flushing of the AdvFS log.
- Fixes an error that can cause a multivolume domain to report ENO\_MORE\_BLKs when some volumes still have free storage.
- Fixes a condition that causes a kernel memory fault.
- Fixes a condition that causes a system hang that occurs when the rename command is called with "." as the target. This patch also reinforces other rename argument restrictions.
- Fixes an RBMT metadata inconsistency that prevents a file system from being mounted.
- Changes a system panic resulting from a kernel memory fault in imm\_remove\_page\_map into a domain panic.
- Fixes a condition that can cause the invalidation of mmap dirty pages before they can be flushed to disk.
- Improves the performance of systems performing heavy file I/O.
- Corrects a potential deadlock hang between a truncate system call and a read system call on a clustered system.
- 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.
- Fixes erroneous logic to ensure that a domain panic on the cluster\_root results in a regular panic for the cluster node on which the domain panic occurs.
- Corrects a condition that causes an "mcs\_lock: no queue entries available" domain panic.
- Corrects a problem in which the cp -p command will not copy DMAPI-managed region information.
- Helps to avoid a kernel panic (kernel memory fault) during a vfixer test.
- Fixes the cause of a system crash when running with lockmode=4. This correction avoids the following types of panics that result from quotactl requests on AdvFS filesets:
 

```
mf in dyn_hash_remove
lock_terminate: lock held
```
- Eliminates the "neg refCnt panic" in AdvFS.
- Fixes a potential panic when using the rmfset command.
- Fixes an infinite loop hang that occurs under special circumstances.
- Fixes an E\_TOO\_MANY\_ACCESSORS error that can occur when deleting a clone fileset.
- Fixes a kernel memory leak that occurs when vfast is in use.
- Fixes a cluster hang where one node tries to get a DIO token and another node tries to start a transaction while a third thread is waiting for the clu\_clonextnt\_lk lock.
- Fixes a potential deadlock hang between a truncate system call and a read system call on a clustered system.
- Prevents a vfast thread from using too much CPU when scanning the AdvFS sbm.
- Provides corresponding memory frees to various mallocs in AdvFS.
- Corrects an infinite looping condition in a vfast thread.
- Corrects a problem in which I/O error codes were not always propagated correctly when AdvFS directIO was used without AIO.
- Fixes a situation that occurs on a full file system in which a write using directIO via AIO may report the incorrect number of bytes written.
- Fixes an unaligned access panic in insert\_seq().
- Fixes a check for an invalid lookup operation through the . tags interface and prevents an AdvFS domain panic.

- Fixes a problem in which a read past the last page of the BMT occurs.
- Fixes the vfast utility error "vfast: cannot get frag list; 14 - Bad address."
- Fixes a "u\_map\_delete failed while deallocating map" error.
- Extends synchronization during directIO writes to include the storage allocation phase.
- Fixes a kernel memory fault that occurs while reading a file with a data inconsistency.
- Fixes a kernel memory fault panic that occurs when recovering an AdvFS domain which was originally crashed for an unrelated reason during a rmfset clone fileset.
- Adds a missing sanity check into AdvFS log recovery code.
- Corrects a condition in which a system would panic due to a stale vdIndex found when writing to a file in an AdvFS domain.
- Fixes a rare race between vfast and mount/unmount in which vfast must open the file differently depending on whether or not the fileset is mounted or not. This fix synchronizes vfast's open with mount/unmount.
- Prevents a node in a cluster from hanging at boot time.
- Prevents a vfast thread from using too much of a CPU when scanning the AdvFS SBM.
- Fixes an AdvFS panic that occurs when deleting an original file and a clone file simultaneously.
- Synchronizes clonefset with read/write paths to force clonefset to complete in a determinate and timely fashion.
- Fixes a problem in which in certain cases the NFS server does not update the access time on the files it serves.
- Fixes a clone data inconsistency that occurs as a result of a remote write to the original cloned file.

#### Patch 27004.00

##### *OSFADVFSDAEMON540*

- Improves the /sbin/init.d/advfsd startup script to allow the user to control the boot time invocation of the advfsd daemon.

#### Patch 27007.00

##### *OSFATMBIN540*

- Corrects a problem that causes a kernel memory fault panic in the event\_queue\_insert() routine on systems using ATM.
- Fixes a problem of stale arp in ATM Elan connectivity.

#### Patch 27009.00

##### *OSFBASE540*

- Corrects a problem with the adduser command so that duplicate UIDs are allowed.
- Corrects a problem seen with -i option of the pr command with respect to white-space substitution.
- Corrects a problem in which the quotacheck utility may incorrectly report that a disk quota has been exceeded.
- Corrects default values for YESEXPR and NOEXPR defined in the localedef command and libc to get correct return value from nl\_langinfo(YESEXPR) and nl\_langinfo(NOEXPR).
- Updates the audit system to display additional information for the numa\_syscalls and msfs\_syscall system calls.
- Corrects a problem with the rm command in which the command fails due to excessive depth when the path name is longer than PATH\_MAX.
- Corrects a problem with the repquota command that causes it to fail with "out of memory for quotause structure."

- Fixes a mailsetup failure if the host name is "unix."
- Corrects a problem in which crash dumps to Fibre Channel swap devices do not always succeed.
- Fixes the following problems with the `acl_set_file()` function:  
`acl_set_file()` fails and returns `errno = 22`  
`acl_set_file()` does not fail if file does not exist
- Corrects a problem with the arena memory allocator. The problem stems from a conflict between `libc` and `libnuma` on handling errors from the `numa_syscalls` system call. For example, an `mfs` request that fails due to insufficient memory should report back a failure message but instead triggers a `segfault` in the caller.
- Resolves a problem that could cause `rexec()` to hang.
- Provides a switch that allows users to specify a port number for `mountd`.
- Fixes `getaddrinfo` so it works properly when IPv6 is not configured.
- Fixes an issue of `sendmail` registration with PSM
- Fixes a problem found during invalid options passed to `mount` while updating UFS file system
- Enables `uudecode` to take care of both absolute and symbolic mode while decoding a file.
- Corrects the message catalogue for `/usr/sbin/auditd`.
- Corrects a problem with the `cp` command by which it continues to be POSIX compliant when the environment variable `STDS_FLAG` is set to the value `ALL`.
- Provides a symbolic representation of the table `syscall` numeric option when using the `audit_tool` in "brief" mode.
- Corrects a problem with the `ex`, and `vi` commands so they use the POSIX compliant shell as the default command line interpreter when the `SHELL` environment variable is set to `NULL`.
- Implements the FC-HBA standard and additional FC-HBA API functionality.
- Fixes an issue with the KZPCC backplane RAID adapter device driver (I2O) that causes its logical disk drives to be identified as SCSI devices.
- Fixes the `chmod` command to ensure that `umask` is taken into consideration when the `who(ugoa)` field is not specified.
- Fixes cryptic error messages generated by the `newfs` command on a volume 1T-512 Bytes.
- Fixes the `sulogin` program to work in a cluster.
- Fixes the `find` command to ignore "--" if it is given as the first argument.
- Provides type checking in `EvmVarGet` wrapper functions.
- Fixes a problem in SIA by resetting the mechanism's context `pkgind` on successful return of `(set|end)*ent` calls.
- Provides enhancements to the `AdvFS` `rmvol` utility, allowing multiple volumes to be removed with one command.
- Fixes a POSIX standard violation in the `strfmon()` function by padding the preceding and following spaces of the returned positive value strings to make an equal length between positive and negative values.
- Fixes a security issue with the C library function `getnameinfo()`.
- Fixes a problem whereby, in a specific circumstance, a memory fault could occur when creating SCS threads.
- Fixes `sdiff` to not show nulls when the `-w` option is used with more than 198 characters.
- Corrects the `audit_tool` output with delimited fields and the output for the "net" and "host" fields.
- Fixes the `/usr/bin/which` command to not display an error message when the environment variable `SHELL` is not set.
- Corrects a problem with the "rcmgr set" command.

- Fixes a POSIX standard violation in the `strfmon()` function by causing preceding and following spaces to be padded to the return value to make an equal length between positive and negative values.
- Fixes the `getnameinfo` routine to display an IPv4 address instead of an IPv4 mapped IPv6 address when the BIND mapping does not exist.
- Fixes issues with "disklabel -e" that occur when extending a disk partition currently in use and open.
- Allows the `fsdb` utility to work on a file system image contained in a regular file, as well as a device. Previously, `fsdb` could analyze only special files, not regular files.
- Corrects a memory leak problem associated with the `clu_get_info` API seen on stand-alone systems.
- Corrects potential security vulnerabilities in the `gzip` program. These vulnerabilities could be exploited by a remote unauthorized user to execute arbitrary code or cause a Denial of Service (DoS).
- Fixes a POSIX standard violation in the `wcstod()` function. Previously, an incorrect pointer was set to the `endptr` parameter of `wcstod()` for cases where no conversion was made.
- Fixes the `awk` command to ensure that the `-f` option exhibits the right behavior with numeric strings.
- Modifies the `ex` command to return 1 as the exit status when a read-only option with write fails.
- Changes the return value of `swprintf()` so it returns the correct value if `swprintf()` detects an invalid wide-character.
- Fixes the `awk` command to ensure that operands that begin with a numeric character or "=" are considered as valid pathnames.
- Updates the time zone data files in `/etc/zoneinfo/` to include the most recent changes from the latest PD time zone source kit (`tzdata2005n`).
- Fixes a problem whereby the `/proc` file system could remain open after a debugging application exits.
- Corrects a potential security vulnerability issue that could result in unauthorized privileged access or denial of service (DoS)
- Implements the FC-HBA standard
- Modifies the `cksum` command to report the correct file size in the `xpg4` environment when the file is greater than 4GB.
- Fixes the problem of `syslogd` dying intermittently when sent a `SIGHUP` to reconfigure.
- Fixes the problem of a core dump occurring while running the `fsck` command on a fake UFS file system.
- Improves space utilization in small MFS file systems by decreasing metadata overhead.
- Fixes a problem with the `mkcdsl` command in which the `-a` option misbehaves when used on a context-dependent symbolic link (CDSL).
- Fixes RPC timeout error messages in the `ypwhich -m` command.
- Corrects a problem in which `SysMan` and `SysMan Station` are not functional after installing Java 1.3.1-4 or higher.
- Prevents `addvol` from adding invalid disks into a domain.
- Reduces cluster files system I/O in Enhanced Security.
- Corrects a security issue in which `rsh` and other `rcmds` incorrectly report `ESUCCESS` when the remote side of a connection terminates before fully establishing a connection.
- Updates `sysconfig` to use the cluster interconnect, thereby allowing for a greater SSI collaboration. This fix will help with changing variables on hung systems, single user systems, and normal running systems.
- Improves cryptic warning message during a mount of a dirty file system. The new message includes a suggestion to run the `fsck` command for a dirty UFS.



- Adds support in cleanPR for HSG60/HSV100/MSA/HP-XP enclosures.
- Corrects a panic while creating or extending a large UFS file system.
- Fixes cron job scheduling for Daylight Savings time events.
- 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.
- Enables the -undo\_switch option for the versw command, which is required to delete patches that contain version switches.
- Corrects a problem in which cp and cat produce different file sizes when reading from a tape device.
- Adds SCSI reserve and release support to the mt program to assist open SAN tape management.
- Updates the environmental monitoring daemon, envmond, to ensure the correct EVM events are being sent at the correct time.
- Allows the auditing of login and su events based in part on the contents of user profiles (for Enhanced Security), the prevailing auditing characteristics of the originating process, and the system-wide audit mask. Previously, only the system audit mask was referenced.
- Fixes many small problems with the dsfmgr command.
- Corrects several potential security vulnerabilities where, under certain circumstances, system integrity may be compromised. These may be in the form of improper file access.
- Adds an event to indicate that the soft or hard error count has changed on the device identified in the event.
- Fixes the search algorithm to differentiate between priviled and non-priviled UIDs, and to allow regular expressions in string searches.
- Corrects a problem found wherein the rmtmpfiles script would leave empty directories in /var/tmp at system startup.
- Fixes a problem that occurs while encoding \$@ in the Bourne shell.
- Eliminates the warning message "Using an array as a reference is deprecated" when running /usr/sbin/siacfg and during system boot on systems using Perl 5.8.0 and higher.
- Corrects a potential floating point error in threaded applications.
- Allows the fuser command to display the reference flag, which indicates the type of reference made; for example, open, closed, unlinked, or mmapped.
- Fixes a problem with csh that occurs when using a tilde (~) operation in directing standard input and standard output of a command to a file in a home directory.
- Updates ddr.mod to support new hardware (NHD-6) devices.
- Corrects a problem in which crontab removes its entries and the vi editor truncates an existing file when a file system is full.
- Corrects the way RPC-based servers handle ill-formed TCP connections.
- Allows mount options that take a value to be correctly processed on a cluster.
- Corrects the tar program to properly handle unusual directory specifications.
- Prevents segmentation faults when a malformed argument vector is passed to sia\_ses\_init.
- Fixes several potential security vulnerabilities where, under certain circumstances, system integrity may be compromised. These 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. This may be in the form of improper file access.
- Corrects a problem in which a core dump occurs when using csh from the Japanese locale.
- Corrects a potential security vulnerability that could result in a Denial of Service (DoS). This potential vulnerability may be in the form of local and remote security domain risks.

(SSRT2384 rpc — Severity - High)

- Corrects a potential security vulnerability in which the Home Directory and login shell attributes for a user account are not supplied to the audit daemon for authentication failures.
- Fixes a problem in XTI caused by a blocked mutex lock in which a thread attempting to send an abortive disconnect hangs.
- Installs DECthreads V3.20-029c.
- Fixes a problem with floating point data inconsistencies in threaded applications.
- Corrects possible dead lock with the `./isl/log` and `./usr/sbin/log` commands.
- Provides the correct labels for mach events to the audit subsystem.
- Corrects the `find -ls` command to display the correct number of blocks.
- 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.
- Corrects the `./usr/sbin/dirclean` utility from attempting to remove the AdvFS `.tags` directory or the `quota.group` and `quota.user` files.
- Fixes an extended regular expression problem where the interval expression `{m,n}` is handled incorrectly.
- Corrects several potential security vulnerabilities where, under certain circumstances, system integrity may be compromised. These may be in the form of improper file access.
- Corrects a 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 memory leaks caused by certain type of scripts that called an infinite loop.
- Fixes a `ksh` problem related to cleaning the process when a terminal is abruptly stopped.
- Corrects the behavior of `ln -sf` to address the issue caused when a symbolic link points to a nonexisting file.
- Corrects the exit status of `sed` when the disk is full.
- Corrects a problem in which the return value of `unlink()` call was not checked when two threads were trying to move a file to two different destinations. Although one of the threads could `unlink()` the source file, no relevant error message was displayed.
- Fixes a problem 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 that have an initially negative `tm_isdst` value.
- Fixes a linker error that occurs when the `ld -update_registry /dev/null` is specified.
- Fixes various problems in the `libc` functions `getdate()`, `strptime()`, `callrpc()`, `strncasecmp()` and `fork()`.
- Fixes a problem in the `libnuma` function `nacreate()` and the system header `<sgtty.h`.
- Causes `sh` to print the correct message when enhanced core file naming is on.
- Fixes a problem in which attempts by the runtime loader (`./sbin/loader`) to free a null pointer are in error.
- Corrects the behavior of the `more` command when nonexisting file and a nonempty file with a long file name are both specified.
- Causes `./usr/opt/ultrix/usr/bin/make` to properly check dependencies on archive libraries.
- Fixes various problems in the `dbx` and object file tools: `dbx`, `ostrip`, `strip`, `mcs`, `dis`, `cord`, `file`, and `stdump`.
- Causes the `grep` command allow blank lines in the pattern file and to not hang when executed with the `-w` and `-f` options.
- Removes compiler warnings addressing outside of array bounds.

- Addresses compiler warnings caused by calling a function with too few arguments.
- Adds informative messages during a rolling upgrade when a problem is encountered with the merging of the .login file.
- Corrects vmstat to display correct free page counts a on NUMA systems.
- Adds a -M command option to newfs to allow permissions of an MFS root directory to be specified when it is first created.
- Adds EVM notification support for UFS file systems.
- Corrects the find -links, -size, -i, -inum behavior with respect to the + operations. Find + operations will match greater than, rather than greater than or equal to.
- Addresses a performance issue of rm -r with large directories.
- Eliminates compiler warnings in ksh.
- Corrects the improper scheduling of cron jobs related to months not having 31 days.
- Makes start up scripts in /sbin/init.d world readable.
- 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.
- Installs version V2.1-120 of /usr/lib/libots3.a and /usr/shlib/libots3.so, which fixes a problem where long-running OpenMP applications might overflow an internal libots3 counter, resulting in a breakdown of thread synchronization.
- Corrects a problem with the merging of the termcap file during a rolling upgrade.
- Allows white space in header field and in multiple headers with the ps option -o.
- Makes the following changes to the tar, pax, and cpio commands:
  - The tar command now checks and reports any write errors.
  - The tar, pax, and cpio commands can unalter the ctime of input files upon creation of an archive and display a warning message if unable to preserve the time of input files.
  - Corrects the tar o option behavior.
  - Corrects the pax -l option to create hard links properly.
  - Corrects the cpio -o option to not ruin extended UID file ownership.
  - Fixes how long file names are handled in tar.
  - Fixes pax to handle ACL on directories properly.
  - Corrects tar to properly handle unusual directory specifications.
  - Modifies the tar utility so it correctly restores directory permissions when extracted using -p option.
  - Fixes the tar and pax utilities so they correctly restore the file mode when extracted using appropriate options.
- Fixes /usr/bin/cut to correctly handle incomplete lines.
- Causes the rescheduling of certain default cleanup cron jobs so that they will not get skipped during a time change to DST.
- Fixes /usr/bin/which to take path information from environment rather ~/.cshrc if it is invoked from other than the C shell.
- Eliminates compiler warnings in mkdir.
- Corrects a problem in which performing a sort on a large database using numerous keys fails during the consolidation phase of the temporary files.
- Fixes a typo in mkcdsl.
- Updates the NIS start-up script to correctly start NIS on the cluster alias.
- Fixes a problem with bcheckrc that occurs when it is run multiple times.
- Fixes a problem with non-U.S. USB keyboards used in non-U.S. locales in which the keyboards are treated as U.S. keyboards by the operating system.
- Corrects a problem in which sh uses a high amount of CPU time.

- Eliminates compiler warnings in ln.
- Eliminates compiler warnings in ksh.
- Enhances the cron command to perform extensive logging.
- Fixes following problems in sh:
  - A service denial problem that occurs when a quoted here doc script is executed.
  - A problem with handling ELF files.
  - A condition in which the shell variable \$- does not hold the -C set option when it is turned on.
  - A condition that causes the printing of broken characters when the type built-in utility of sh is invoked in the Japanese locale.
- Fixes a one-byte gap/hole in the maximum file size in the tar command before an extended header record is used (8589934591 (octal 7777777777)).
- Fixes a problem in which nonsense characters are appended to the audit information output of an execve event in brief mode.
- Installs DECthreads V3.20-033, which addresses the possibility of floating point errors in threaded programs.
- Corrects a problem in which some networking applications, especially X.25 and X.29, stopped working as expected because of interactions with security-related fixes and how the fstat() function behaves on their sockets.
- Corrects a potential security vulnerability which may result in non-privileged users gaining unauthorized access to files or privileged access on the system.
- Fixes an sh problem that occurs when executing command substitution.
- Fixes a fatal assertion error reported by pixie, hiprof, third spike, cord, uprofile and odump object file tools for some executable files linked at optimization level 2 (-O2) or greater.
- Corrects a condition in which the system start-up script /sbin/rc2.d/S19security could cause an unintended change in the system security configuration, which could occur when /usr/bin/perl has been removed.
- Corrects a condition in which the mv command fails when operated on running binaries.
- Corrects a condition in which the mv command does not allow a file to be renamed even though directory permissions allow it.
- Fixes a problem encountered with the Bourne shell when a file name with trailing slash (/) is used as an argument to the command.
- Allows the mv command to parse the pathname correctly when a source directory ends in trailing slash (/).
- Corrects the default UID displayed by the adduser command when UID\_MAX exists in the /etc/passwd file and helps prevent the duplication of UIDs.
- Corrects a potential security vulnerability in sendmail 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 or remote security domain risk.
- Corrects a condition that causes NIS clients to fail to connect to non-Tru64 NIS servers that support only the V2 NIS protocol.
- Fixes a situation in which the header fields displayed by the output of the ps command are not aligned properly.
- Corrects a problem with the class scheduler under the following conditions:
  1. The class scheduler is started then stopped.
  2. System owned semaphores are removed with the ipcrm -s command
  3. The class scheduler fails to restart with the error "class\_open: allocate or access semaphore Invalid argument."
- Fixes a situation, in which the mkpasswd command dumps core when executed by a nonprivileged user.

- Fixes a situation in which awk does not process input files specified in the BEGIN section.
- Corrects a problem with the cdvd command that causes two minor problems with its output.
- Fixes a problem that occurs when linking kernel object files from an archive library. The problem was the linker always adds four `__exc_*` symbol references to the final linked image.
- Corrects a problem that can cause binlogd to dump core when parsing its remote host authorization file (`/etc/binlog.auth`) with greater than 513 characters.
- Removes the 250 variable limit for `/usr/bin/env`.
- Fixes a problem of race condition in `rm` command in which two threads can successfully delete a file simultaneously.
- Makes it possible for the `vi` command to read a text file containing non-ASCII single-byte characters (for example, accents) and to read single-byte locales correctly.
- Corrects a condition that causes the FTP daemon to dump core when the client sends an out of order ADAT command.
- Causes EVM to check for matching braces within EVM configuration files.
- Fixes a buffer overflow problem in `/usr/bin/write`.
- Adds a CDSL from `/var/adm/binlog.saved` to `/var/cluster/members/{memb}/adm/binlog.saved`.
- 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, sending the message "Impossible to execute `/sbin/sh`."
- Removes the race security vulnerability in the `find` utility.
- Changes the `sort` command to give exit value 1 for all the error messages, in compliance with existing specifications.
- Prevents the `sort` command from dumping core when more than 50 sort keys are used
- Provides the feedback facility for `dd` in long copies.
- Adds support to the `evminfo` command to check for syntax errors in the EVM authorization file.
- Fixes a problem with the UFS file systems that occurs after using the `extendfs` command.
- Fixes two problems in the dynamic runtime loader that can cause an application to crash.
- Fixes a problem that causes a segmentation fault when `dbx` is analyzing a Fortran program.
- Fixes various problems regarding regular expressions in multibyte locales.
- Corrects a possible hang problem with complex regular expression.
- Corrects a potential security vulnerability that can result in unauthorized Privileged Access or a Denial of Service (DoS). This may be in the form of local and remote security domain risks.
- Fixes a possible data inconsistency that can occur when copying files across DMAPI-enabled file set using the `cp` command.
- Addresses a problem with `sh` while using the `ulimit` built-in command in displaying hard and soft resource setting values when `-H` (hard) and `-S` (soft) resource limits options are specified.
- 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.
- 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.
- Causes the appropriate exit status to be returned when a disk is full.
- Lets `awk` accept input records of a length up to 5,119 bytes.
- Fixes problem with the `vi` internal command when used inside a macro to repeat the last search for a pattern.
- Corrects the last command to correctly report logout times when several sessions are active.
- Adds support for the `file` command to recognize ELF and HP-UX file formats.

- Corrects a condition in which a user would receive a protocol-handshake error when a high priority binlog event is reported via an email message. (The problem is caused by EVM's use of DECEvent to translate the binlog event, which requires the HOME environment variable to be set.)
- Corrects the behavior of munlockall in the realtime library (librt).
- Corrects a condition that causes a panic resulting from a kernel memory fault in access\_invalidate.
- Corrects a problem in mount or domain activation after a panic, where a fileset (domain) cannot be mounted without running fixfdmn.
- Corrects a setld security issue in relation to SSRT 3471.
- Causes environmental monitoring on Alpha Server ES47/ES80/GS1280 systems to be turned on by default.
- Corrects problem of excessive swapping resulting from mfs file system creation.
- Fixes a problem with mkcdsl not carrying the sticky bit through.
- Fixes a problem to allow the class scheduler to handle processor sets (partitions) with an ID greater than 100 when attempting to set the partition (using the setup subcommand). The error message generated by the problem is "invalid partition specified."
- Fixes a ksh memory fault problem that occurs when logging in.
- Corrects a problem on systems running Enhanced Security in which the command edauth -R refuses to write user-profile entries to the root partition.
- Corrects problems with name resolution when an error is encountered during the processing of the local host files.
- Fixes a fatal error in /usr/bin/spike.
- Allows the Event Manager daemon, evmd, to stop listening on its default TCP port 619. This capability is not available for clustered systems.
- Corrects a problem that occurs when using C1crypt for password encryption on Enhanced Security systems in which users are unable to change their passwords and see the passwd command warning "Password not changed: failed to write protected password entry."
- Corrects the output of the vmstat command with per-RAD kernel usage data.
- Fixes problems such as segmentation faults when strxfrm() function runs on the French locales.
- Fixes setld failures that generate the message "Kerberos credentials not found."
- Improves null partition checking code.
- Fixes a librt memory leak that may occur when multiple message queue files are opened and then closed (the memory would be recovered when the process terminates).
- Changes the use of the configuration file /etc/svc.conf to /etc/nsswitch.conf to allow netgroup data to be provided from LDAP, rather than only from NIS.
- Corrects a potential security vulnerability that can result in a local Denial of Service (DoS).
- Addresses the rarely seen class scheduler error "class\_daemon: class\_open: Database created with default configuration Bad file number."
- Fixes the pattern substitution problem in ksh with respect to the patterns \${parameter#pattern} and \${parameter%pattern}.
- Adds per-binary big page controls to complement the system-wide tunable attributes.
- Updates the audit\_tool usage message.
- Fixes a performance problem in the libc mktime() routine.
- Fixes a condition that causes segmentation faults when the nm -a command is used to dump symbol information in object files compiled by the C++ compiler.
- Fixes the dircmp to display both output columns.
- Fixes a condition that causes cfgmgr to continue running after the remote is gone.

- Adds a -R switch to MFS to allocate memory on one RAD. The default MFS behavior is to stripe memory across all RADs that have memory; the -R switch may improve MFS performance in many cases.
  - Fixes a problem that prevents access to AutoFS file systems if ACLs are enabled.
  - Makes the cchwtest tool for Common Criteria security evaluation.
  - Fixes the compress command to exit with new status, 3, on an I/O error, as described in the compress reference page.
  - Corrects a problem in which sendmail will not start on some configurations where the sendmail support files are symbolically linked to another location in the support environment.
  - Fixes “at” jobs that fail to execute when LANG and LD\_LIBRARY\_PATH are set.
  - Modifies the dc command to take care of scale while printing numbers in base 2.
  - Causes the tar utility to correctly restores directory permissions when extracted using -p option.
  - Fixes the vi command with respect to the handling of locales and the +[subcommand].
  - Corrects the behavior of the find command with respect to -size option.
  - Corrects the behavior of the ltf command for UIDs and GIDs greater than four digits.
  - Eliminates several compiler warnings.
  - Expands libpset APIs to enable the caller to get processor set information.
  - Modifies the rewind() function to always reposition to the beginning of a file.
  - Modifies the tar and pax utilities to correctly restore the file mode when extracted using appropriate options.
  - Eliminates inefficiencies and apparent login hangs in a TruCluster environment running Enhanced Security.
  - Corrects the handling of shared library version mismatches for the loader's -ignore\_version and -ignore\_all\_versions switches.
  - This modification adds an additional helpful instruction to the instructions given to users after an SIA initialization failure.
  - Corrects a fixso failure that occurs when reporting a new conflict for \_\_istart.
  - Corrects the behavior of the w command, when it is redirected or piped.
  - Corrects a failure in the safe\_open() routine that causes symbolic links given by a relative path from the current working directory to incorrectly give ENOENT errors.
  - Installs DECthreads V3.20-029, which fixes problems that may affect threaded programs.
  - Fixes a problem where a file name is assigned incorrectly in a CDSL case.
  - Installs DECthreads V3.20-049.
  - Corrects a problem in which the evmshow command occasionally encounters print formatting problems when displaying a detailed output of binlog channel events.
  - Corrects an odd, unexpected error message that may be printed by rsh or rlogin commands.
  - Corrects a problem where the telnet command causes unnecessary delays when an IP address is supplied as a command-line argument.
  - Fixes a problem in which the mount command dumps core or gives inappropriate warning messages if a malformed option was given.
  - Corrects a problem in which the mailsetup command does not allow changes to the local user list in a noncluster environment.
  - Corrects a potential security vulnerability that may result in unauthorized Privileged Access or a Denial of Service (DoS). This may be in the form of local and remote security domain risks.
- (SSRT2384 rpc — Severity - High)
- Adds support for NEW\_OPEN\_MAX\_SYSTEM (64K) file descriptors to libaio.
  - Corrects a problem in which a DNS resolver routines never time out if interrupted by signals.

- Corrects a problem in which the pop3d command dumps core when SSO (single sign-on) is configured.
- Corrects a problem in which long-running programs using Enhanced Security interfaces (such as getespwnam) directly grow in memory use over time.
- Provides an RFC3542 compliant implementation of IPv6 Advanced API.
- Fixes a ksh core dump problem that occurs when too many files are opened; for example, when executing too many scripts simultaneously.
- Fixes the more command to handle multibyte character properly.
- Fixes a nonconcurrency issue for multithreaded applications calling popen() and certain "FILE \*" routines such as fread().
- Corrects the /sbin/sulogin utility to read values correctly from the /etc/rc.config.common and /etc/rc.config.site files.
- Fixes a deadlock condition in multithreaded applications that call fork() and other libc callback routines such as exit handlers, \_\_fini\_\* routines.
- Eliminate spurious "dln\_tsl\_set: abort: resource name changed" messages and the associated authentication failures when /var/tcb/files is no longer in the /var file system.
- Installs DECthreads V3.20-049a, which fixes a condition that may cause some threaded applications to hang.
- Fixes a problem with the find command that occurs when it is used with -follow option on symbolic links to a directory.
- Adds support for latest version of the bcm card.
- Corrects a condition in which Multiple "Sorry" messages are issued by the su command when multiple SIA mechanisms are in use (as when LDAP is configured for user accounts).
- Updates Mobile IPv6 code to be compliant with the latest RFC.
- Modifies the /usr/sbin/versw\_enable\_delete script, which is used to enable the deletion of version switched patches.
- Enhances the audit\_tool to enable a user to specify a path to archived audit logs that the audit\_tool will follow for all logs instead of the path recorded in the audit\_log\_change events (which by default is /var/audit).
- Fixes an error reported by dbx when the process being debugged terminates.
- Corrects a problem in which the evmshow command hangs when processing an invalid evm.buf file.
- Corrects a problem in which the system hangs at boot due to evmd processing an invalid evm.buf file.
- Fixes a vold threads problem in which LSM devices may be incorrectly recorded as disk clones when LSM starts up.
- Provides a new LSM EVM event that is posted when LSM is processing clusterwide plex detaches.
- Corrects a problem in which xdmcp terminal failures are not logged within the terminal control database for Enhanced Security.
- Corrects a problem in which the Mail program would occasionally generate duplicate file names when invoking the editor.
- Corrects the default answer for cleanPR clean.
- Corrects the behavior of the svr4 -t option of the df command in a cluster environment.
- Allows V5 auditd servers to communicate with V4 auditd clients.
- Fixes a problem with csh that may cause a "Missing }" error.
- Fixes problems with Enhanced Security user accounts that reference a template.
- Corrects a problem in which stdin, stdout and stderr are closed and cfg\_connect fails to connect to cfgmgr.



- Provides an enhanced version of the rm command that fixes a problem with VMS NFS mounts in which many files are not deleted because of shuffling effects in the file system.
- Fixes a timestamp problem with binlogd
- Fixes an issue with setld/pax.
- Fixes a memory fault or segmentation fault problem that occurs when a root user tries to change a NIS user's password on the NIS master that is using BASE security.
- Corrects a problem with the Bourne shell not displaying the trailing slash at the end of a file name that is passed as an argument to the script that displays the argument passed.
- Fixes a problem in which software can not be installed due to subsets remaining in an unfinished state of subset loading.
- Corrects a csh globbing problem that occurs while listing files using the asterisk (\*) wild card character.
- Fixes a problem with the sh shell during command substitution.
- Fixes the problem of depord failing when the length of the DEPS field in the control file is greater than 250.
- Fixes quotacheck utility to handle user/group IDs larger than  $2^{31}$ .
- Fixes the auditlogtrim script to properly recognize and delete audit log files older than the trimming date/time specified.
- Ensures that the disklabel -z option (used for zeroing out a disk label) issues a warning if the disk partition is in use when a shell environment variable DISK\_ZERO\_WARN is set
- Corrects an overlapping section error reported by the linker.
- Fixes bad error handling that occurs when trying to install setld via the command line options setld -l <path <subset.
- Corrects a memory leak problem in the ksh shell.
- Corrects bind changes to rc.config changing NIS settings.
- Fixes a problem that results in a hang or crash during the auditing of the net\_tcp\_stray\_packet, net\_udp\_stray\_packet, and net\_tcp\_rejected\_conn network events .
- Fixes and audit object deselection model in which events that modify a target object could be deselected, resulting in no audit record being generated.
- Improves the performance of times(3).
- Corrects incorrect behavior of the df command for automount mount points.
- Corrects a memory leak problem in the ksh shell.
- Enhances the fuser command to provide a cluster-wide query capability. A revised fuser(8) reference page describes this enhancement.
- Fixes a problem with glob() returning an incorrect match when directory permission issues exist.
- Corrects a condition that occurs when a change in system time by the ntp daemon causes a cron job to be executed with a discrepancy equal to the amount of drift in time.
- Corrects the behavior of the tail -f command with respect to FIFO special files.
- Corrects a problem with ksh shell pattern matching.
- Fixes problem that occurs when multiple instances of the sort utility use the same directory path.
- Resolves several DECThreads faults and resolves performance issues with certain Java applications.
- fixes memory leak problem in cron.
- Raises the highest ID handled by edquota and repquota and converts the repquota program to using the GETQUOTA/GETQUOTA64 quotactl's for obtaining data, instead of directly reading from the quota file. This results in more recent data and, in a cluster, consistent data across all nodes. It also enables handling the highest UIDs (up to 4294967294).

- Fixes an awk() argument buffer limitation that can cause patch installations to fail if a large number of subsets (in excess of 500) are loaded on the system.
- Fixes sh to handle a large number of open files.

## Patch 27010.00

### *OSFBIN540*

- Provides support for identifying references to half-open pipes via the fuser utility.
- Improves performance for some multithreaded applications running on AlphaServer GS320 and GS1280 class systems.
- Corrects a scenario in which a spurious wakeup is sent to a process that had interrupted a blocked attempt to set a mandatory file lock.
- Provides performance enhancements to the vm\_overflow feature.
- Changes the way page migrations occur on a NUMA system to address poor performance due to excessive paging.
- Corrects an incompatibility between the cpus\_in\_rad and gh\_chunks/rad\_gh\_regions tunable attributes that could result in a boot failure.
- Corrects a problem with pagetable page allocations that could leave a thread waiting indefinitely during a fork operation.
- Corrects "ubc\_wire: hash failed" panics on non-NUMA systems.
- Corrects a "not wired" panic that occurs with System V shared memory and bigpages.
- Reduces scheduling contention when unmapping shared address space that has an extremely high number of mappers.
- Handles reservation conflict errors to address cluster node hang during boot.
- Updates the audit system to display additional information for numa\_syscalls and msfs\_syscall system calls.
- Provides a tunable attribute to allow the reserving of a percentage of vnodes for root use.
- Fixes an issue in the VM subsystem in which a page that is not managed by VM is incorrectly identified as being managed by VM.
- Fixes a problem in UFS superblock update logic that could result in a failure to report errors encountered when writing cylinder group summary blocks.
- Provides a new cluster-specific link aggregation distribution algorithm when using LAG in a LAN cluster.
- Reduces fork/exec overhead on systems with large default stack sizes.
- Fixes a cluster deadlock/hang issue when a new device is discovered.
- Fixes a kernel memory fault panic that occurs during a NetRAIN interface failure.
- Fixes a CPU hang caused by /dev/random.
- Corrects a condition that causes a process hang and system panic when using System V shared memory and asynchronous I/O.
- Corrects a minor tuning issue with netisrthreads on NUMA machines. Previously, four netisr threads were started per RAD, whereas the default number of netisr threads per RAD now equals the number of CPUs per RAD.
- Fixes problems with RPC for configurations with no RAD 0.
- Fixes a problem in which unkillable processes occur when debugging multi-threaded programs with the totalview debugger on a Sierra Cluster.
- Fixes an underlying problem in NFS that can trigger an assert failure in CFS. These conditions can be triggered by a failed NFS mount.
- Resolves a synchronization issue between pageout and exit paths.
- Introduces sysconfig tunable attributes for NFS/RPC
- Fixes multiple panics and application hangs seen when interacting with Process Shared POSIX 1003.1c objects.

- Fixes the handling of NFS requests that have an erroneous file handle length field.
- Corrects the output of wired pages and gh regions found in the -P option of the vmstat command when run on a NUMA system.
- Improves the performance and chances of success of the Tru64 kernel malloc routine, especially when used by drivers which use the M\_NOWAIT option.
- Fixes an invalid kernel memory fault panic in the code responsible for the generation of random numbers.
- Corrects a defect in the audit subsystem that causes it to fail to record inode information on closed file descriptors.
- Corrects a problem that occurs after an upgrade from Version 5.1B patch kit 3 to Version 5.1B-3 (Version 5.1B-3) in which the system may fail to boot with the following panic:  
KMF - invalid memory read access from kernel mode
- Fixes a race within NFS over TCP when connection reaches idle timeout.
- Improves the performance of Tru64's MACH implementation and also enhances stability while under heavy load.
- Corrects several potential security vulnerabilities in TCP/IP, including ICMP. These exploits could result in a remote Denial of Service (DoS) from predictable Initial Sequence Numbers (ISNs), network throughput reduction for TCP connections, or the reset of TCP connections.
- Corrects the following potential security vulnerabilities:  
SSRT4743, SSRT4884 - TCP/IP ICMP (Severity - High)
- Fixes a panic condition resulting from DVDFS using a deallocated vnode.
- Fixes a panic condition in the NFS client code resulting from the mishandling of unaligned data on clustered systems.
- Corrects a kernel memory fault that results from the de-referencing of a null processor pointer encountered when auditing a network event after having sent a SIGKILL to the auditd daemon.
- Forces the use of a non-cluster interconnect address in audit records.
- Fixes a kernel memory fault panic in IPv6, and corrects improper or leaked reference counts with IPv6 route entries.
- Fixes a panic that occurs when using the TCP SACK option.
- Fixes a kernel memory fault panic that occurs when the vm tunable anon\_rss\_enforce is set.
- Fixes a socket kernel memory leak.
- Fixes a race condition in the UBC.
- Corrects tracking of controlling terminal reference to session structure.
- Corrects a problem that causes the transmission of duplicate FIN packets, which could result in a stuck connection.
- Allows systems with multiple paths to a large number of devices to boot faster.
- Fixes an issue with the NFS client async daemon that can occur in specific NUMA configurations. These configurations would have RAD numbers that are not part of the partition but within the range of valid RAD numbers for that partition.
- Fixes an fsync (NFS/UFS) operation that failed to always flush all dirty pages.
- Fixes a "lock\_done: lock not currently owned" panic that occurs early in a boot.
- Resolves lock management issues within the UBC that can lead to "mcs\_unlock: current lock not found" and "mcs\_lock: time limit exceeded" panics.
- Removes NFS warning message about malaligned RPC messages.
- Makes the poll() function compliant with revised UNIX98 standards.
- Fixes a hang caused by the Async I/O subsystem.
- Corrects a problem with the mount command in which issuing a mount -l command lists the property of a dual mounted fileset incorrectly as nodual.

- Fixes a problem that causes some TCP connections to reset when sending large amounts of data (more than 2GB) to a very efficient receiving host.
- Corrects the incompatibility of the waitpid() system call with revised UNIX98 standards.
- Fixes an internal kernel declaration for variable maxuthreads that resolves the following types of kernel failures when the sysconfig value for max\_threads\_per\_user in the proc subsystem is set to a value larger than 2147483647:  
thread\_create() failed for pid # : maxuthreads (=1) exceeded for uid 201  
thread\_create() failed for pid 1479950: maxuthreads (=1) exceeded for uid 15
- Fixes a panic resulting from CDFS using a deallocated vnode.
- Corrects a problem in which the tmt\_walk\_list() function in nifftmt.c mallocs memory for the tmt\_state\_info array, but fails to free memory when thread terminates.
- Reduces the default value of IPFRAGTTL from 60 to 18 to avoid reassembly problems.
- Modifies the setluid system call to handle the license correctly when a non-root user switches to different non-root user using login command.
- Corrects a problem in which elevated load averages were reported on NUMA class systems.
- Fixes resource leaks seen after a device file is revoked.
- Makes possible a sticky connection feature for cluster alias.
- Updates sysconfig to use the cluster interconnect to allow for greater SSI collaboration, which will help with changing variables on hung systems, single user systems, and normal running systems.
- Allows the UFS attribute delay\_wbuffers to be tuned using sysconfig.
- Allows the packet filter variables pfil\_loopback and pfil\_physaddr to be tuned using sysconfig.
- Increases the default value of ipqmaxlen (IP input queue) to 2048.
- Provides an option to enable cluster NFS clients to use a nonprivileged TCP port to check to see if a remote NFS server is up.
- Corrects a potential issue with NFS version 3 memory mapped files that can lead to a system panic.
- Allows the stat system call to correctly report the st\_blocks for dvdrom files.
- Corrects a condition that causes a panic while creating or extending large UFS file system.
- Corrects a problem in which under certain load conditions shared memory usage can lead to an inconsistency that results in a "u\_ssm\_oop\_deallocate: reference count mismatch" panic.
- Corrects a potential hang on exit from applications utilizing /dev/poll.
- Fixes a condition in the kernel whereby an incorrect internal status can be returned from mpsleep(), thereby causing potentially incorrect behavior.
- Fixes problems in tcp\_output that cause connections to hang when window fills and fixes a potential loss of data when a connection is closed.
- Corrects a potential loss of data a connection is closed.
- Corrects inappropriate TCP probe timeouts associated with case.
- Adds support for CPU offline on GS1280 systems (required for Capacity on Demand).
- Corrects a problem with the keepalive mechanism that causes TCP connections to disconnect unexpectedly.
- Corrects a problem in which changing memory protection results in a kernel memory fault panic.
- Fixes a condition that can cause a panic in the kernel during interconnect operations involving configuration requests.
- Corrects a problem in which an invalid core file may be generated following abnormal program termination.
- Fixes mount/umount failures and panics in FMS, UFS FDFS.

- Fixes a situation in which mmap memory locked with mlockall() using the MCL\_FUTURE flag does not become wired automatically.
- Fixes a problem in which fuser is unable to report on all referenced resources when attempting to identify reasons for unmount failures.
- Improves the process exit procedure for processes that have had the nice command used on them.
- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file access.
- Fixes multiple defects in AutoFS user space and kernel code.
- Corrects performance issues when accessing a file with direct I/O enabled.
- Fixes a condition that causes a panic when appending to a file.
- Eliminates a false directory lookup warning message generated by an incorrect comparison caused by mismatched file ID variable types.
- Improves client caching performance.
- Prevents the loss of a single system image for an NFS file system mounted from a cluster when there are certain problems communicating with the external NFS server.
- Fixes a problem with audit data not being displayed by the audit tool.
- Corrects problems with file object selection and deselection and directories.
- Addresses NUMA performance issues associated with auditing.
- Introduces type checking of attributes when registering components with the hardware manager.
- Adds IEEE 802.1Q (VLAN) support.
- Fixes a panic within the two-level scheduling subsystem.
- Fixes a process hang condition.
- Fixes a "thread\_block: simple lock held" panic.
- Fixes an occasional panic that can be seen when reading from a process using Granularity Hints with the procfs command.
- Corrects a problem that causes a system to panic during a particular machine check.
- Corrects 3D client hangs when using the Radeon graphics card.
- Fixes a condition that causes a system to hang when using Open3D over the AGP bus on an AlphaServer GS1280.
- Protects against "get\_color\_bucket: empty buckets!" panics and "kernel memory fault" failures on systems with mixed cache parameters.
- Fixes a kernel memory fault in shadowvnode() caused by a null vnode pointer.
- Fixes insmntque() to conform to proper locking when removing and adding a vnode to the mount vlist.
- Fixes a condition that causes an excessive FIDS\_LOCK contention when large numbers of files are using system-based file locking.
- Fixes an issue encountered in configurations in which the primary processor is not the first processor within a RAD.
- Fixes a condition that causes the panic "u\_seg\_vop\_remove: seg not found" in vop.
- Fixes a problem in which a duplicate IP address might be configured on the system or an IP address might be configured with an incorrect netmask.
- Extension of UFS file systems via the mount command can effectively disable use of the file system. Additionally, on some LSM based systems a panic can occur after a file system extension has been completed.
- Fixed a problem in which a process waiting on a semaphore does not get woken up.
- Fixes a problem that causes the panic "Bigpage Assertion Failed."
- Allows the size of the NFS server's duplicate request cache to be adjusted as needed.
- Corrects a locking problem with NFS running over UFS.

- Fixes a problem in which a Tru64 UNIX NFS client panics when it receives a null entry as a response to a readdirplus request from an NFS server.
- Fixes a problem in which a Tru64 UNIX NFS server panics as a result of receiving illegal file access mode from an NFS client.
- Corrects a rounding error for the vm attribute vm\_bigpg\_thresh.
- Corrects the handling of bad pages when big pages are enabled.
- Fixes the cause of "page mapped" panics when using mmap calls with dev/mem to access free big pages.
- Increases the TCP window increased from 96 KB to 500 KB to improve performance.
- Causes the netisr thread to dynamically estimate the reply size the socket buffer and subsequently reserve the space in it.
- Adds a new timeout check to notice when data has not been acknowledged in 30-50 seconds and copy those buffers, thereby allowing the UBC to free up those mbufs and not tie them up.
- Adds support for CPU indictment on the AlphaServer GS1280 platform.
- Fixes a problem in which gh\_min\_seg\_size can not be set below 8M.
- Fixes a problem with cluster failovers of UFS filesets.
- Fixes a panic resulting from race condition in the MFS (memory file system) over CFS (cluster file system).
- Corrects a problem in which an ARP request for a permanent ARP entry is ignored and a connection cannot be made from the remote system.
- Fixes a high lock contention for str\_to\_lock, a STREAMS attribute.
- Adds a configurable tuning parameter to the STREAMS subsystem for ES47/ES80/GS1280 platforms.
- Fixes a flaw in the NFS server that can cause it to crash upon reception of malformed input.
- Fixes a kernel memory fault in u\_seg\_global\_destroy.
- Corrects a kernel memory fault that can happen when running applications that use the Cray Intra-Node Shared Memory library.
- Prevents a potential process (not system) hang seen when a system comes under heavy memory load with monolithic memory use. (Gigabyte-scale single objects.)
- Prevents a kernel memory fault when running with protection on the 128-byte bucket — an action that should only be undertaken as directed by HP customer support personnel.
- Addresses a situation in which a taso-compiled binary is unable to allocate more memory after performing a series of mmap calls.
- Corrects a problem that causes hwmgr to dump core when performing environmental testing and using hwmgr to verify that a particular sensor's status would change from OK to Fault.
- Provides the correct labels to the audit subsystem for mach events.
- Fixes a memory management fault and panic in UFS.
- Addresses two problems with the NFS server:
  - A potential crash during a concurrent read and truncate operation on an AdvFS file.
  - A potential crash with malformed or malicious READDIR[PLUS] version 3 RPCs.
- Improves ufs\_invalidate() handling of fractional pages.
- Corrects the cause of an "ialloc: dup alloc" panic.
- Increases the character limit in file property lists from 245 to 255.
- Improves I/O performance by reducing kernel locking overhead.
- Addresses system problems that can occur when the system is under heavy I/O load or low memory conditions.
- 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 a kernel memory fault panic that occurs when the audit daemon is set to periodically dump the kernel audit buffers to the audit log file (auditd -d freq).
- Corrects conditions that cause "blkfree: freeing free block" and "blkfree: freeing free frag" panics.
- Allows systems configured to use NTP to keep accurate time regardless of whether the NTP daemon is running.
- Corrects a condition that can cause a panic in audit\_rec\_build when auditing execve with the exec\_argp or exec\_envp audit style enabled.
- Fixes a problem in which a device file such as /dev/console can become inaccessible, returning the error "Bad File Number."
- Fixes a system panic in the ubc\_page\_stealer routine.
- Causes the correct error message for the freezefs -q command to be displayed on a non-AdvFS file system.
- Adds comment to reserve 0x10000000 and 0x20000000 for AutoFS flags.
- Prevents issues in the DCE/DFS file system when pages are being flushed as part of a vnode. This patch is required for AlphaServer SC.
- Fixes a problem where an I/O error (EIO) can occasionally be returned after a page fault.
- Adds a lock to the initialization of a private client to avoid hung file system threads when the MVFS ClearCase file system is in use.
- Corrects a potential security vulnerability that may result in denial of service. This may be in the form of local and remote security domain risks.

(SSRT2384 rpc — Severity - High)

- Corrects the NFS server's handling of files open for direct I/O.
- Fixes a problem on a cluster NFS client where a hard-mounted NFS file system can incur ETIMEDOUT errors.
- Fixes a problem in the kernel network subsystem that causes a kernel memory fault panic in the m\_adj() routine.
- Fixes an internal problem in the kernel's AdvFS, UFS, and NFS file systems where extended attributes with extremely long names, greater than 247 characters, could not be set on files. The new limit is 254 + a null string terminator.
- Fixes a problem when the kernel incorrectly closes a socket that causes Sybase 1613 errors.
- Adds support for IPV6\_UNICAST\_HOPS socket option on raw sockets.
- Corrects c shell a problem in which a multithreaded process forks a single threaded process and leaves data in the proc structure that could cause problems.
- Improves performance for removing or truncating large files on UFS file systems.
- Fixes a problem that occurs when the Tru64 UNIX TCP layer prematurely closes a slow, but good connection with TCP reset.
- Fixes a problem in which the quot -v command sometimes returns wrong quota information on a UFS partition.
- Fixes a system panic that can generate the panic string "mcs\_lock: lock already owned by CPU" or "thread\_block: simple lock owned."
- Fixes a system panic that generates the panic string "pg\_nwriters going negative."
- Changes an rws write lock in the VMAC lookup routine to an rws read lock for better SMP scaling.
- Corrects a kernel memory fault in the table syscall.
- Corrects memory striping when using big pages.
- Fixes a problem where the system can panic with a kernel memory fault in simple\_lock() being called from fuser().

- Resolves kernel memory faults in the TCP/IP subsystem.
- Fixes a problem where threads can hang mallocing memory.
- Increases the default values for `udp_ttl` and `tcp_ttl` to 128 hops.
- Fixes a condition that causes the "rdg: unwiring" panic.
- Increases the default limit of DLI packets to 16 KB and makes the limit tunable. This corrects a problem in which attempts to send packets larger than 5000 bytes (jumbo packets) can fail.
- Improves the fragment gathering mechanism to boost performance.
- Corrects a problem in which the `auditd -d` command (which flushes the kernel audit buffers) could cause audit data inconsistencies on a multi-CPU machine on systems generating a heavy volume of audit events.
- Provides a base system routine for use by the cluster code to determine if a particular mount has an NFS exported file or directory on it.
- Corrects a problem in which a user process cannot be interrupted, which in some instances can be utilized in a denial of service attack.
- Corrects a kernel memory fault caused by uninitialized or incorrect parameters being passed to the `setsockopt` system calls.
- Fixes a condition that causes a kernel memory fault panic in the IP multicast loopback code.
- Corrects a problem in which multi-CPU sometimes livelock while processing incoming network traffic. In some cases the live lock can result in a cluster event timeout panic.
- Resolves internet protocol conformance issues and fixes a problem with sending multicast datagrams.
- Fixes a condition that can cause crashes from within the `pshared` subsystem
- Fixes several IPMI-related problems, including the following:
  - Erroneous fields in 686 OS-detected environmental machine check logout frame
  - Unusually large number of 686 sensor timeouts with heavy system load
  - IPMI always reporting -48v sensors as broken, seen as "redundant power supply failed" messages
  - An IPMI memory leak
- Corrects problems in UFS `extendfs` functionality that cause file system metadata inconsistencies.
- Corrects the default parameter for `physio_max_coalescing` to 8K.
- Addresses an issue on large systems in which kernel threads might not be executed for extended periods of time.
- Fixes two small logic errors with the NFS version 3 client that result in unintended, though correct, behavior.
- Fixes a problem seen with the TAHI IPv6 conformance test, specifically Test 4 for the IPv6 specification.
- Resolves a problem that results in multiple cluster members crashing with kernel memory fault in `rfs_find_fsid()`.
- Corrects a potential system hang when an error occurs while updating special file or named pipe access times on an NFS client.
- Fixes an AdvFS asynchronous direct I/O problem that can cause a thread to hang.
- Fixes a problem in which a truncated AdvFS file erroneously zeros data for the remaining leading segment of the file.
- Fixes a minor problem in the IPv6 subsystem that causes an extra message to be sent upon startup.
- 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.
- Improves performance for CFS filesets mounted with the `server_only` option.



- Allows tasks with just one SCS thread to be migrated in the same manner as single threaded processes in NUMA environments.
- Addresses a problem in IPv6 subsystem that causes a system to panic.
- Adds support for AlphaServer ES47/ES80/GS1280 platforms to allow the use of processors having different speeds in the same chassis.
- Corrects the problem of the statfs() function returning EINVAL when operating on an fattach() function clone streams device.
- Addresses a scaling issue seen on large multiprocessor systems in dealing with the class scheduling subsystem.
- Fixes a problem in gh\_chunks allocation on some configurations.
- Improves file caching performance for large files on NUMA systems.
- Fixes panic: simple\_lock\_terminate: lock busy.
- Fixes cluster crashing when VMAC is enabled
- Fixes an underlying problem in the NFS client that can lead to a panic on a single system or an assertion failure panic on a cluster.
- Fixes the number of cylinders defined in the last cylinder group of a UFS file system that has been extended.
- Fixes various problems in the bcm driver for DEGXA Gigabit Ethernet adapters.
- Removes the 32-bit block address restriction in SpecFS read and write.
- Removes excess data stored in the auditing of the swapctl syscall.
- Adjusts the automatic computation of the size of the NFS server's duplicate request cache to accommodate only active RADs on a multi-RAD system.
- Increases ability to audit memory wire and unwire operations.
- Resolves a problem affecting IPv6 sockets where a packet transmission might fail when the destination address is not specified.
- Fixes a kernel memory fault panic while running the IDRIS benchmark.
- Corrects a kernel memory fault caused by configuring an IPv6 address on a legacy network interface card.
- Fixes multiple problems with IPv6 advanced API implementation.
- Corrects a problem in which calling aio\_write with a negative size may cause a system panic.
- Corrects a potential security vulnerability reported by SSRT2323.
- Fixes a leak of kernel address space.
- Corrects a problem in which a read or write operation to a changer device creates an unkillable process.
- Corrects the panic "trap: invalid memory read access from kernel mode in procsf\_psinfo()."
- Corrects a condition in which big page memory allocations on NUMA systems may allocate remote memory too soon.
- Fixes a problem in which clua.mod does not handle TCP RST messages appropriately.
- Updates ICMP code to add redirect timeout support and nonmodifiable route table entries.
- Addresses performance issues with the select() function.
- Accelerates the booting of memory GS systems.
- Adds pfil\_loopback and pfil\_physaddr kernel flags for controlling packetfilter written packets.
- Corrects a "simple\_lock: time limit exceeded" issue that can occur in the CAM I/O subsystem when MCS locking is disabled.
- Provides support for the dynamically loadable packet filter.
- Fixes IP multicast packets to work with loopback traffic
- Improves NFS client performance on NUMA systems by doing a better job distributing data, locks, and nfsiod threads across Resource Affinity Domains (RADs).

- Enables suitable monitoring of individual NetRAIN interfaces
- Provides a tunable attribute that gives system administrators the ability to adjust the weight that the kernel places on the NUMA locality for forks on AlphaServer GS80, GS160, GS320, and GS1280 systems.
- 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.
- Adds per-binary big page controls to complement the system-wide tunable attributes.
- Enhances HP-XP array controller support and possible future new tape device support.
- Fixes a panic condition in AIO.
- Fixes a "simple lock: time limit exceeded" panic.
- Fixes a problem that prevents access to the AutoFS file systems if ACLs are enabled.
- Improves NFS client performance on NUMA systems by doing a better job distributing data, locks, and nfsiod threads across Resource Affinity Domains (RADs).
- Addresses the dispatching of NFS server requests for the public file handle on cluster members and multi-RAD systems.
- Fixes how the NFS client handles full NFS version 3 64-bit file IDs.
- Resolves a problem that prevents the viewing of files created by third-party software on some CD-ROM media.
- Resolves a problem of not being able to view files on some CD-ROM media that is created by third-party software and corrects the erroneous reporting of success when attempting to write beyond the file size limit using synchronized I/O and the calculation of `_PC_FILESIZEBITS`, which is used by the operating system for pathconf file characteristics.
- Fixes a race condition in the kernel AIO code that can panic the system with either a kernel memory fault or a duplicate malloc free.
- Corrects a potential panic with large memory processes using System V shared Memory at process exit.
- 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).
- Fixes a memory leak in the NFS server encountered when it receives malformed packets.
- Fixes an AutoFS panic during an unmount operation when AutoFS tries to remove a directory.
- 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.
- Corrects a problem in which `table()` calls were not correctly getting process arguments.
- Improves the scaling of IP reassembly code on large SMP machines. NFS servers are especially susceptible when a large number of clients attempt to write at the same time.
- Implements buffer cache page checksum caching for NFS client pages.
- Fixes a problem that can cause a system crash when an NFS server exports files on a third-party file system (that is, one not built into Tru64 UNIX).
- Fixes a rare kernel panic that occurs during the handling of a clock tick when class scheduling is used.
- Corrects a potential process hang that occurs while exiting a system that has dynamically powered off a processor.
- Corrects a system panic when running with big pages enabled.
- Addresses problems that occur when taking noninteractive core dumps using the `coredump` command.
- Corrects potential hangs of applications using the pshared subsystem that can occur as a result of a thread failing to wake up after its condition variable timer had expired.

- Fixes two potential problems in the NFS V3 client in which unstable writes can remain uncommitted when they should have been committed to stable storage.
- Removes erroneous "No B-cache detected" messages from certain configurations.
- Prevents a potential panic that generates the message "memory\_test=partial" or "memory\_test=none."
- Allows the stat system call to correctly report the st\_blocks on a CD-ROM file.
- Improves performance for applications with large in-core data sets
- Corrects a system panic caused by stack growth.
- Adds the ability to join more than 20 IP multicast groups on a given socket.
- Fixes a multiple process hang (which cannot be terminated with Ctrl/c) that can occur if process A attempts to attach process B to another RAD while at the same time, process B attempts to attach process A to another RAD (a classic deadlock).
- Changes the implementation of the NFS server's duplicate request cache from a statically allocated monolithic entity to a dynamically allocated entity.
- Corrects a kernel memory fault panic under certain heavy system loads when using the /proc file system to debug processes.
- Corrects a problem in which using PRSABORT in the /proc file system does not correctly abort a system call in the process being debugged.
- Fixes a "simple\_lock: time limit exceeded" panic involving the vm\_object.ob\_lock lock.
- Fixes a rare case of a thread blocking when waiting for memory.
- Fixes a problem in which some IP fragments of NFS over UDP read replies may be sent to the wrong MAC address.
- Permits the setting of the setuid, setgid and sticky bits on NFS served files.
- Fixes kernel memory faults that can occur when packetfilter support is used incorrectly as a result of the following:
  - A misconfiguration of the packetfilters system configuration parameter in the net subsystem.
  - The dynamic load of the packetfilter subsystem on a system that has packetfilter statically built into the kernel.
- Fixes a condition in which the fork() function returns EAGAIN incorrectly on a process using gh\_chunks.
- Corrects a problem in which interface Unicast packet counters seem to go backwards when retrieved via SNMP.
- Improve the performance of systems that are performing heavy file I/O.
- Corrects several minor problems with the IPv6 subsystem related to the Neighbor Discovery specification.
- Provides an RFC3542 compliant implementation of IPv6 Advanced API.
- Fixes a panic condition caused by a problem in the swapping subsystem.
- Allows the niffconfig command to exceed 10 interfaces.
- Corrects a potential kernel memory fault.
- Fixes a problem of a system hang that occurs when the system may be swapping
- Fixes overall system instability caused by the pshare subsystem
- Speeds up the discovery of a healthy cluster interconnect interface by helping NetRAIN failover to a redundant healthy interface before the (Internode Communication Services) begins to time out on the keepalives, thereby preventing the cluster from crashing.
- Fixes certain panics conditions by increasing the timeout for CPU onlines.
- Updates Mobile IPv6 code to be compliant with the latest RFC.
- Corrects the handling of a condition that can lead to a panic under heavy system load with high swap activity.

- Corrects a problem in which the `cp -p` command will not copy DMAPI-managed region information.
- Addresses an issue in which the default value for kernel flag `mobileip6_enabled` is 1
- Corrects the cause of a panic on a cluster system when renaming a root specified by the `chroot` command.
- Fixes a kernel memory leak that occurs during routing table updates, which over time, can cause memory shortages for machines attached to a constantly changing network, as in the case of mobile IP.
- Fixes the cause of a hang/panic that occurs soon after a console `printf` returns the message "NFS over TCP client xxx not reading replies, continuing."
- Corrects a kernel memory fault in `exit()` where `uu_curinfo` is `nil`.
- Fixes `MIP6_CACHE_READ_TO_WRITE_LOCK()` to not call `lock_read_to_write()` when in `lockmode=0` to avoid a panic.
- Modifies asynchronous I/O to prevent it from releasing file region locks.
- Fixes "issig recursion" panics.
- Fixes a simple lock timeout panic in the `ubc_invalidate` routine.
- Corrects a condition in which the `fork()` function can lead to a "vm\_pg\_free: page mapped" panic when big pages are enabled.
- Corrects a problem with HSZ events that cause I/O errors rather than doing the appropriate retries so the I/O can be successful.
- Corrects problem that prevents proper XP RaidManager operation.
- Corrects a Routing Header Type 2 length miscalculation.
- Fixes a problem in which some reads on `/dev/random` wait until the entropy pool is completely full while other reads of `/dev/random` and `/dev/urandom` do not.
- Fixes a problem on a VMAC-enabled cluster wherein the member that acts as the Proxy-ARP-Master for the default cluster alias address sometimes incorrectly uses the VMAC address created for it as the source MAC address.
- Fixes the cause of a "u\_anon\_unlock\_page: anon already unlocked" panic.
- Fixes a problem in which the NFS version 3 server can reset the `setuid` and `setgid` bits of a file or directory after attempting to comply with the specified attributes in the request.
- Addresses a performance issue seen in certain applications that allocate and exercise a memory segment close to the system's `bcache` size.
- Prevents aio-related slowdowns.
- Corrects an internal error so that UBC invalidate operations can execute correctly.
- Fixes the cause of an "mcs\_unlock: lock not currently owned" panic.
- Addresses problems with reading third-party file systems exported by the NFS server.
- Fixes crashes seen on clustered systems.
- Addresses the panic "pmap\_enter\_bigpage: attempt to map big page in kernel map" that is sometimes seen on large AlphaServer GS1280 configurations.
- Fixes an AutoFS problem, where indirect-mapped key directories are not removed for failed mounts.
- Fixes a potential deadlock hang between a `truncate` system call and a `read` system call on a clustered system.
- Enables AutoFS mount-on paths to have the correct maximum length when AutoFS files are mapped directly.
- Fixes a condition that causes the `evmwatch` or `evmshow` commands to display incorrect IPv6 addresses.
- Corrects a problem with handling of null symbolic links.
- Exports the creation time of a process via its `/proc` entry's `ctime` attribute.

- Fixes a problem in which questionable behavior by non-TRU64 NFS clients results in NFS over TCP server threads hanging.
- Fixes a problem related to the deletion of binding cache entries in mobile IPv6, which affects multi-RAD systems like the AlphaServer GS1280.
- Addresses issues that may be seen as panics or hangs having to do with the UBC.
- Introduces the `vm_overflow` tunable.
- Provides enhancements for NetRAIN operations.
- Corrects race in streams I/O completion and timeout logic.
- Adds the ability to decouple RPC client retransmission activity from UDP or interface transmission completion handling.
- Fixes a condition that causes a connection timeout error when using socket keepalive options with small values on sockets over a localhost connection.
- Corrects the cause of the a panic: "panic: Hashing that is not marked VPP\_REPL."
- Corrects a problem in the kernel dli module that causes system hangs and crashes on multi-CPU systems. The crashes occur with the panic string "lock\_terminate: lock held".
- Fixes a logic error in the virtual memory macros used for accessing the memory descriptor, based on physical address.
- Corrects the panic "trap: invalid memory read access from kernel mode" in pshared code.
- Corrects a problem with the randomization of initial TCP sequence numbers.
- Allows multiple retries to mount the root file system and the ability to adjust the retry period.
- Enhances the range of generic sysconfig parameters and eliminates unnecessary sysconfig warning messages.
- Corrects several problems related to use of the `ifconfig [interface] IPv6 up` command, where interface is either cluster interconnect, `ics0`, or configured v6-in-v6 tunnel.
- Fixes an NFS client kernel loop that usually results in a lock time limit exceeded panic on partitioned GS1280 class systems in partitions where CPUs have different numbers of neighbors. That usually requires a partition of more than four CPUs.
- Ensures that the socket send buffer size will be twice the maximum segment size if the mss is 1460 bytes.
- Fixes a problem that causes a hang or crash when auditing the network events "net\_tcp\_stray\_packet," "net\_udp\_stray\_packet," and "net\_tcp\_rejected\_conn."
- Corrects a problem with the audit object deselection mode in which it was possible for events that modify a target object to be deselected, resulting in no audit record being generated.
- Corrects for a leak of CONTROL mbufs.
- Fixes a problem in which a threaded program tries to wake up a thread that is already running.
- Enhances the `fuser` command to provide a cluster-wide query capability. A revised `fuser(8)` reference page describes this enhancement.
- Fixes a problem in which the `traceroute` command will timeout when probing the cluster alias address from a cluster member.
- Corrects a memory leak against the 128-byte bucket when `rad_gh_regions` are configured and `nshmget()` is used.
- Adds support for `ifconfig [inet6] delete [abort]`, for IPv6 connections.
- Corrects a problem in which the `netstat -s` command does not display the correct global statistics on multi-CPU systems if lower-numbered CPU slots are empty.
- Corrects a memory leak in IPv6 transmission code.
- Corrects a condition that results in a failure to release file region locks (NFS client only).
- Corrects a panic "vm\_pg\_free: page on o/h list."
- Fixes a process deadlock that occurs when `rename` is called with "." as the target.

- Resolves a potential system hang in the kernel virtual memory subsystem when running the ClearCase V5.0 Multiversion File System (MVFS).
- Fixes a problem that prevents 6to4 and autotunnels from being disabled properly when the corresponding IPv4 address is deleted.
- Fix to send maximum sized packets on FDDI networks when pmtu\_enabled=0.
- Makes Tru64 UNIX V5.1B compatible with V5.1A in regard to the handling of the relationship between a controlling terminal and the "session" structure.
- Fixes the fastpath send of RPC modules to correct conditions in which it ignores the VMAC setting and it excludes source and destination UDP port numbers when LAG is enabled.
- Fixes a problem in which the NFS server, in certain cases, does not update the access time on the files it serves.
- Enables IPv6 MTU to be set up to 9000 bytes on gigabit Ethernet.
- Fixes memory troller panic that can occur when booting in low-memory conditions.
- Addresses two potential causes of various ics timeout panics, such as "ics\_unable\_to\_make\_progress: input tread" and "ics\_unable\_to\_make\_progress: netisrs stalled."
- Enables allows AutoFS mount-on paths to have the correct maximum length when AutoFS files are mapped directly.
- Fixes a race condition in VM map-entry fault handling that can result in poor performance and/or panics.
- Fixes a SierraCluster "bigpage assert failure" panic.
- Corrects the panic:"kernel memory fault" in umc\_page\_release() due to non-initialization of variable hpp.
- Fixes a problem that can cause the system panic "cmn\_err: ce\_panic: ics\_unable\_to\_make\_progress: netisrs stalled."
- Fixes a problem that causes multiple TCP sockets to be created and then left in the CLOSE\_WAIT state on an NFS client in a TruCluster when the NFS server is unavailable.
- Provides improvements to the TCP Selective Acknowledgment (SACK) option.
- Corrects an issue in which tape devices with no ddr entry would not have compression correctly enabled or disabled as appropriate to the device special file that was used.
- Corrects a problem that causes data transfers to be slow on socket connections when using the MSG\_WAITALL flag in a recv() socket call.
- Corrects a condition to avoid the copying of multicast packets to the networking stack when packetfilter is enabled and pfilt\_copymulti=0; default is 1, meaning always copy multicast packets.
- Fixes NFS client hangs on systems with no associated memory for RADs.
- Modifies the kernel to correct a problem that could under certain conditions cause the following:
  - User space programs to experience unexpected "interrupt" results (EINTR) returned by system calls such as exit() or wait3().
  - Korn shell scripts to report that programs were terminated with a "Signal 64" when in fact the program ran normally and no error was present.
  - Programs that check the exit status of children to see abnormal results.
- Fixes a problem of ambiguous connection rejection caused by socket database search in a multi-RAD system.
- Corrects several problems in file region locking that could result in dropped file region locks or in a process hang.
- Fixes an "m\_copydata offset" panic seen on TruClusters and NUMA systems serving NFS file systems.
- Fixes a race condition in the UBC subsystem that could result in a kernel memory fault.

- Prevents a race condition in the code dealing with kernel address space data structures.
- Improves NFS to help prevent an NFS client from creating more than one TCP connection for a mount point and to prevent a system panic if the problem occurs.
- Fixes a lockmode 4 only panic due to a lock ordering problem.
- Resolves a problem that could result in freeing active memory buffers.
- Prevents the growth of a stack object when code invokes an mmap system call.
- Corrects a kernel memory fault panic clock\_tick().
- Fixes the following panics:  
page is not primary page  
mcs\_unlock: current lock not found
- Corrects a memory leak caused by a race condition that can lead to the incorrect decrementing of reference counts of vm\_anon objects, which can cause orphaned pages.
- Fixes an "mcs\_lock: lock already owned by cpu" panic.
- Fixes a problem that can trigger a "big pages assert" panic under the following condition:
  - Big Pages are enabled.
  - Large argument lists are enabled.
  - An application such as a shell tries to pass a large list of arguments to a child process, such as a command.
- Fixes the following panics:  
mcs\_unlock: lock not currently owned  
the zombie walks, the sequel
- Changes I/O retry processing when a command time out is encountered by causing the retry code to attempt to take advantage of other active paths during retries if command time outs are encountered.
- Fixes a rare lock timeout on some AlphaServer ES47/ES80/GS1280 systems with heavy binary errorlog activity.
- Fixes a "vl\_unwire: page is not wired" panic when the system is configured with gh\_chunks or rad\_gh\_regions enabled and new\_wire\_method = 1.
- Retires the "new\_wire\_method" sysconfig tunable parameter. See "new\_wire\_method Tunable Attribute Retired".
- Corrects a potential security vulnerability in the Transmission Control Protocol (TCP) that could be remotely exploitable, resulting in denial of service (DoS).  
SSRT4696 - TCP (Severity - High)
- Corrects a potential security vulnerability that could result in a Denial of Service (DoS). This potential vulnerability may be in the form of local and remote security domain risks.
- Fixes a panic in ubc\_bigpage\_release().
- Fixes a problem that causes the audgen system call to overwrite memory beyond the end of the audit buffer.

## Patch 27011.00

### *OSFBINCOM540*

- Fixes a rare kernel memory fault problem that occurs during an ES80 boot.
- Provides a new cluster-specific link aggregation distribution algorithm when using LAG in a LAN cluster.
- Allows for control ports to be deleted using the hwmgr command.
- Fixes a potential deadlock hang between a migration and a flush on a file.
- Addresses a potential I/O performance bottleneck in which the tape driver may select the same path for all tape drives in a system containing multiple HBAs and multiple tapes. The

tape driver can now assign different paths to different tape drives to improve tape I/O performance.

- Fixes problems that happen when volume expansion (`mount -u -o extend`) races with other code.
- Corrects a C++ compilation error when including `conf.h`.
- Makes the `poll()` function compliant with revised UNIX98 standards.
- Revises the `btcreate` utility to overcome the 32MB firmware limitation on SAS kernel size.
- Reduces the default value of `IPFRAGTTL` from 60 to 18 to avoid reassembly problems.
- Provides an option that allows cluster NFS clients to use a nonprivileged TCP port to see if a remote NFS server is up.
- Allows the `hwmgr redirect scsi` command to work with lockmode 4.
- Corrects a problem in which the heavy use of the cluster alias on large SMP and NUMA machines results in excessive context switching and CPU load.
- Enables SmartArray 5300 controller hardware events to be logged to the `binary.errlog` during boot time. This is useful in diagnosing logical volume state change and physical drive hotswaps that can occur while the system is not booted.
- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file access.
- Adds SCSI reserve and release support to `mt` to assist open SAN tape management
- Fixes a problem with audit data not being displayed by the audit tool.
- Corrects problems with file object selection and deselection and directories.
- Corrects NUMA performance issues associated with auditing.
- Adds support for IEEE 802.1Q (VLAN).
- Adds support for IEEE 802.1Q (VLAN) (`DE50x`, `lan_common.h`).
- Adds support to get live status information for air movers and power supplies on AlphaServer GS1280 systems and to log intrusion packets to the error log.
- Adds support for CPU indictment on AlphaServer GS1280 platforms.
- Fixes a panic condition in which the operating system erroneously reboots instead of halting and fails to take a crash dump.
- Fixes numerous issues in the driver for DEGXA Gigabit Ethernet adapters, including the DS25 on board 10/100/1000 port.
- Adds defensive programming to `stat.h` to prevent it from getting confused if one of its internal temporary `#defines` is defined before `stat.h` is processed.
- Corrects a problem in which the `btextract` utility was not preventing the advanced mode of restore for a system with LSM setup.
- Removes files created under `/usr/lib/sabt` during the running of the `btcreate` utility on file systems with LSM. Previously, these files were copied to the tape, and would be restored as if they were archived by `btcreate`.
- Addresses problems with the `mksas` utility in which false warning messages are generated and the user-specified temporary directory can be erroneously removed.
- 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.
- Fixes the IDE/ATAPI driver.
- 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 a process hang in `ubc_common_lookup`.
- Corrects the NFS server's handling of files open for direct I/O.



- Prevents problems in the USB subsystem that include memory leaks, data inconsistencies, and USB device configuration problems.
- Fixes problems in the USB driver layer that include potentially minor performance degradations if using USB devices and occasional USB device failures.
- Corrects a problem in which a temporary file is left behind during the creation of bootable tape with the UFS file system.
- Fixes an underlying problem in the NFS client that can lead to a panic on a single system or an assertion failure panic on a cluster.
- Fixes various problems in the bcm driver for DEGXA Gigabit Ethernet adapters.
- Defines new error log entry types for system error logging on certain 21364-based systems.
- Increases the ability to audit memory wire and unwire operations.
- Corrects a problem in which a read or write operation to a changer device creates an unkillable process.
- Fixes a problem in which clua.mod does not handle TCP RST messages appropriately.
- Increases BOOTP error timeout values for RIS kernel installations.
- Accelerates the boot time on large memory GS systems.
- Provides support for dynamically loadable packet filters.
- Distributes the latest .h files for the mcutil program and the event manager.
- Adds function prototypes for the NUMA APIs `cpu_get_current()` and `cpu_get_rad()` to `sys/cpuset.h`.
- Adds the function prototype `numa_query_pid()` to `numa.h`.
- Changes the use of `/etc/svc.conf` to `/etc/nsswitch.conf` to allow netgroup data to be provided from LDAP in addition to NIS.
- Provides support for the Philips USB controller, which is shipped on some AlphaServer GS1280 systems. Without this patch, the Philips USB controller may fail to detect and configure devices connected to it.
- Sets a software limit to prevent serious performance problems that can occur when TCP connections that have an rate limit enforced by a downstream network device overrun the device.
- Corrects a problem in which `/sbin/ddr_config` does not accept values for `ReadyTimeSeconds` larger than 255. The new limit is 86400 seconds (24 hours).
- Fixes problems with NUMA disk statistics.
- Fixes a KMF problem that can occur when some nodes in cluster are rebooted and a device is shared by all the nodes.
- Changes the CAM subsystem message that is printed to the error log on a recovered read error from "bad block number" to "block number."
- Adds per-binary big page controls to complement the system-wide tunable attributes.
- Fixes how the NFS client handles full NFS version 3 64-bit fields.
- 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 with `hwmgr` utility deletes while a SCSI scan is in progress.
- Improves scaling of IP reassembly code on large SMP machines. NFS servers are especially susceptible when a large number of clients attempt to write at the same time.
- Implements buffer cache page checksum caching for NFS client pages.
- Corrects a kernel memory fault caused by configuring a system to be a firewall router without having the `GWSCREEN` subsystem built into the running kernel.
- Improves performance for applications with large in-core data sets.
- Corrects a problem in which the less than (<) and greater than (>) keys on the Norwegian keyboard do not work in single-user mode.

- Adds the ability to join more than 20 IP multicast groups on a given socket.
- Changes the implementation of the NFS server's duplicate request cache from a statically allocated monolithic entity to a dynamically allocated entity.
- Fixes a problem in which some IP fragments of NFS over UDP read replies may be sent to the wrong MAC address.
- Improves the performance of systems that perform heavy file I/O.
- Provides an RFC3542 compliant implementation of IPv6 Advanced API.
- Provides support for latest version of bcm card.
- Updates Mobile IPv6 code to be compliant with the latest RFC.
- Fixes MIP6\_CACHE\_READ\_TO\_WRITE\_LOCK() to not call lock\_read\_to\_write() when in lockmode=0 to avoid a panic.
- Corrects a problem with HSZ events that cause I/O errors rather than doing the appropriate retries so the I/O can be successful.
- Adds read/write ability to XP control ports.
- Corrects a problem with a sensor error that can indicate a false over-temperature condition on DS10/DS10L and TS10 systems.
- Fixes USB problems that prevent some USB keyboards and mice from working.
- Provides enhancements for NetRAIN operation.
- Adds the ability to decouple RPC client retransmission activity from the UDP or interface transmission completion handling.
- Fixes a logic error in the virtual memory macros used for accessing the memory descriptor.
- Allows multiple retries to mount the root file system and the ability to adjust the retry period.
- Fixes a problem in which the operating system receives an environmental machine check packet from the firmware but fails to correctly recognize the sensor that is identified as faulty by the machine check.
- Fixes a problem that prevents keyboard input during interactive booting (SRM boot flag "i") when in graphics mode on GS1280/ES80/ES47 AlphaServers.
- Enhances the fuser command to provide a cluster-wide query capability. A revised fuser(8) reference page describes this enhancement.
- Fixes mass storage disk subsystem handling of specific important events reported by the HSG80 and HSG60 array controllers.
- Provides improvements to the TCP Selective Acknowledgment (SACK) option.
- Provides a workaround for the firmware limitation on SAS kernel size.
- Raises the highest ID handled by edquota and repquota and converts the repquota program to using the GETQUOTA/GETQUOTA64 quotactl's for obtaining data, instead of directly reading from the quota file. This results in more recent data and, in a cluster, consistent data across all nodes. It also enables handling the highest UIDs (up to 4294967294).
- Fixes a problem that corrects the status message which is output during boot operations before the binary.errlog log file is available.
- Fixes a memory troller panic during a boot in low-memory conditions.
- Fixes a rare kernel memory fault that can occur when booting an AlphaServer ES80.
- Changes I/O retry processing when a command time out is encountered by causing the retry code to attempt to take advantage of other active paths during retries if command time outs are encountered.

## Patch 27012.00

### *OSFC2SEC540*

- Fixes problems with `prpasswd` and `rpc.yppasswdd` that can cause these daemons to consume high CPU time when run in a TruCluster Server environment.
- Allows root to log in on the console when Enhanced Security is enabled and `u_numunsuclog` exceeds `u_maxtries`.
- Modifies the `prpasswd` and `rpc.yppasswdd` daemons to properly handle `/var/tcb/files` on a file system from different from `/var`.
- 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 can hang on any member other than the one serving `/var` to CFS.
- Fixes a problem in which group and other read privileges get stripped from `/etc/passwd` when a user switches from enhanced to base security.
- Corrects a problem on systems running Enhanced Security in which the command `edauth -R` refuses to write user-profile entries to the root partition.
- Corrects a problem that occurs when using C1crypt for password encryption on Enhanced Security systems in which users are unable to change their passwords and see the `passwd` command warning "Password not changed: failed to write protected password entry."
- Changes the use of the configuration file `/etc/svc.conf` to `/etc/nsswitch.conf` to allow netgroup data to be provided from LDAP, rather than only from NIS.
- Fixes problems with Enhanced Security user accounts that reference a template.

## Patch 27013.00

### *OSFCDEAPP540*

- Resolves security vulnerabilities within the X PixMap routines used in the IMG library.
- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file access.
- Corrects a problem in which `dtcm` dumps core when the following steps are taken:
  1. Select a date
  2. Choose Browse-Compare Calendars...
  3. Press the Mail... button.
- Fixes the `dtcm` warning message when selecting View → Day in `dtcm`.
- Resolves a potential buffer overflow within the X PixMap routines.

## Patch 27014.00

### *OSFCDEDEV540*

- Fixes several potential security vulnerabilities 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.
- 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 problem that can occur when the screen saver tries to activate on a system that has reached the maximum number of processes allowed per user and the following message is displayed:

An attempt to start a new process on host "hostname" failed

- Corrects a problem in which the application builder core dumps when trying to generate code for menu items with the set-label action type.

## Patch 27015.00

### *OSFCDEDT540*

- Corrects several potential security vulnerabilities where, under certain circumstances, system integrity may be compromised. These may be in the form of improper file access.
- Fixes the message catalog for the CDE application dtprintinfo.
- Fixes a problem in which dtcreate core dumps while editing the icon image.
- Corrects a problem in which setting the value of screen saver and screen lock in dtstyle disables both values.
- Corrects a problem that causes dtsession to not work when its norestore option is enabled.
- Resolves a problem that occurs when displaying the user-specified logo in dtlogin.
- Resolves the incorrect system activity report by the w command with XDMCP.
- Fixes a dtcreate problem that occurs while saving an action file in an NFS-mounted environment.
- Fixes the dtsetup display.
- Sets defaults for dtlogin.
- Changes the quick setup message.
- Sets the proper text for the front panel help page in dtsetup.
- Corrects a potential security vulnerability in the Common Desktop Environment (CDE) software. This potential vulnerability, which may be locally and remotely exploitable, could result in a denial of service (DOS), unauthorized privileged access, or both.  
SSRT4721 - dtlogin - (Severity - High)
- Ensures that the proper error message is set in dtlogin when an unknown display parameter is read from the Xaccess file.

## Patch 27016.00

### *OSFCDEMAIL540*

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

## Patch 27019.00

### *OSFCDEMIN540*

- Resolves a problem that occurs when opening big file using dtfile.
- Fixes several potential security vulnerabilities where, under certain circumstances, system integrity may be compromised. These may be in the form of improper file access.
- Corrects a 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.
- Fixes dtterm problem that causes a core dump when the resource saveLines values is set to 1000s in \$HOME/Dtterm.

- Corrects a problem that can occur when the screen saver tries to activate on a system that has reached the maximum number of processes allowed per user and the following message is displayed:  
An attempt to start a new process on host "hostname" failed
- Corrects a potential security vulnerability in CDE code that may result in unauthorized privileged access. This may be in the form of local and remote security domain risks.  
(SSRT3589 - dtmailpr Severity - High)
- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised when a buffer overflow occurs in the CDE online help. Buffer overflows are sometimes exploited in an attempt to subvert the function of a privileged program and possibly execute commands at the elevated privileges if the program file has the setuid privilege.
- Fixes a dtmail problem that occurs while opening a mail attachment on an NFS-mounted environment.

## Patch 27020.00

### *OSFCLINET540*

- Fixes a problem with the niffconfig command whereby certain characters in the interface name may be ignored.
- Fixes a problem in the /etc/.mrg..inetd.conf merge script that causes customer-specific changes in the /etc/inetd.conf file to be ignored.
- Fixes a problem with inetd -L in which a cluster loops in shutdown -c or rcinet start.
- Upgrades BIND 8 to BIND 9.
- Fixes a problem that occurs when starting inetd on all RADs in which there are holes between the RADs.
- Fixes a potential remotely exploitable Denial of Service (DoS) vulnerability in the File Transfer Protocol server daemon, (ftpd) in which under certain circumstances authorized users could cause an ftp server to become unresponsive.
- Adds a -n option to the ftpd daemon to prevent login delays and time-outs in an environment where host name resolution is sluggish.
- Adds a new table in pm.mib for the pmgrd IoRate Statistics feature.
- Adds the file pmAdvfs.MIB to define AdvFS MIB definitions.
- Allows the optional port argument to the ftp open command to accept port numbers between 32768 and 65535.
- Fixes several potential security vulnerabilities where, under certain circumstances, system integrity may be compromised. These may be in the form of improper file access.
- Corrects a potential security vulnerability that may allow nonprivileged users to gain unauthorized (root) access. This may be in the form of local and remote security domain risks.
- Corrects a potential security vulnerability in BIND 8 code that could result in a local or remotely exploited Denial of Service (DoS).  
(SSRT3653 - BIND v8 — Severity - High)
- Corrects a problem in niffd that results in its memory usage growing over time.
- Fixes a problem in the operation of the IPv6 neighbor discovery daemon where IPv6 addresses are not automatically configured on PPP interfaces.
- Adds support for IEEE 802.1Q (VLAN).
- Fixes a problem that prevents startslip from extracting all the information from the acucap file.

- Fixes a problem in the `/etc/.mrg..protocols` merge script that causes incorrect permissions on the `/etc/protocols` file.
- Corrects the `netstat` and `ifconfig` commands so that when a MAC address is printed, it uses 2 digit hex octets with leading zeros.
- Corrects a potential security vulnerability that may result in a Denial of Service (DoS). This potential vulnerability may be in the form of local and remote security domain risks. (SSRT2384 rpc — Severity - High)
- Corrects a problem in `os_mibs` that results in the swap size and swap used values for the host mib being reported as negative values on some systems.
- Introduces `dumprmt.msg` for remote dump/restore messages. This new message catalog file is used in both `rdump` and `rrestore` programs.
- Corrects a problem in which Solaris `rcp` commands fail against Tru64 UNIX V5.1B servers with the message "rcp: lost connection."
- Enables the `tip` command to log into the member-specific log file. The path for the `aculog` file has been corrected and given appropriate permissions.
- 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.
- Makes start up scripts in `/sbin/init.d` world readable.
- Provides support for monitoring the disk I/O rate using the `pmgrd` daemon.
- Adds support for monitoring AdvFS statistics using Performance Manager agent, `pmgrd`.
- Fixes the `mkcdsl` command and updates the NIS start-up script to correctly start NIS on the cluster alias.
- Resolves a problem in which the cluster interconnect route is inappropriately advertised.
- Corrects a problem in which the `nissetup` command leaves `/etc/group` with an incorrect mode of 600 after removing NIS.
- Corrects a problem in which NIS clients may fail to connect to non-Tru64 NIS servers that only support the V2 NIS protocol.
- Corrects a problem that causes the FTP daemon to dump core when the client sends an out of order ADAT command.
- Adds support the `gated aliases-next-hop` option, which provides a preference for the selection of a next-hop address when multiple addresses exist on the interface.
- Corrects the `/sbin/init.d/route` script to ensure that routes get flushed properly.
- Corrects a problem in which if a static or loopback host route is added via a routing socket, `gated` deletes the route upon a route aging time-out. In cluster alias environment, this problem causes all TCP/UDP packets destined to the cluster alias address to be mishandled.
- Adds support to the `ifconfig` application for the IPv6 command line argument `ip6reachabletime`.
- Fixes a condition that causes an `ftpd` file transfer failure when a file exceeds 2 GB.
- Fixes an `ftp` client accounting error that occurs when transferring files larger than 4 GB.
- Fixes library dependencies to allow `ifconfig` to be run in single user mode.
- Corrects a problem that occurs when using MD5 authentication with Version 2 RIP.
- Corrects a condition that causes a hang of the IPv6 routing daemon.
- 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 the `gated` daemon or `gated` control utility (`gdc`) incorrectly access temporary files.

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of mouted daemon incorrectly creating and accessing temporary files.
- Adds a required call to yp\_unbind() after getnameinfo() in order to close all ports when using telnet on some machines.
- Provides comment handling capability in the route initialization script for parsing /etc/routes file.
- Fixes an incorrect diagnostic message in the traceroute command.
- Modifies ifconfig to include an RCSid for patch tracking purposes.
- Updates ICMP code to add redirect timeout support and nonmodifiable route table entries.
- Fixes setld failures with "Kerberos credentials not found" messages.
- Changes the use of the configuration file /etc/svc.conf to /etc/nsswitch.conf to allow netgroup data to be provided from LDAP, rather than only from NIS.
- 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.
- Eliminates the message "usage: hostid [hexnum or internet address]" from the /usr/sbin/niscluster script while using the -d option.
- Resolves a problem that occurs when forking inetd children processes on NUMA-based AlphaServer platforms.
- Addresses a problem with a missing file that is required for the pmgrd daemon's iorate metrics collection feature to work.
- Modifies NIS slaves running in a cluster to use CAA to help assign a crontab entry to update NIS maps.
- Fixes the ftpd SIZE subcommand to conform to RFC 959.
- Resolves a problem that can cause rsh processes hang.
- Corrects a problem where the telnet command causes unnecessary delays when an IP address is supplied as a command-line argument.
- Corrects a problem with the nslookup utility that results in an unexpected termination in certain situations.
- Fixes the behavior of rpc.rquotad in a cluster.
- Allows rcp to correctly interpret file names that contains rcp special characters such as the at ( @ ) sign and colon ( : ).
- Fixes a problem with the IPv6 neighbor discovery daemon, which under certain circumstances, can cause bad information to be written to a DNS database causing failures on subsequent database reloads.
- Corrects several minor problems with IPv6 subsystem related to neighbor discovery specification.
- Corrects a problem with routed not detecting all the interfaces on system that contains more than 100 interfaces.
- Updates Mobile IPv6 code to be compliant with the latest RFC.
- Provides support for monitoring disk I/O rate using the pmgrd daemon.
- Resolves a problem in which a hang occurs after the "User logged in" message is displayed during an ftp session to a UCX host as a user with no password.
- Resolves a problem in which gated can produce a core after port 616 is scanned by a port scanner.
- Fixes a problem with FTP byte/hash count when the file size exceeds 2 GB.
- Modifies niffconfig to display appropriate error messages.
- Corrects a problem in which ftp reget works incorrectly when restarting at 2 GB or larger.
- Provides enhancements for NetRAIN operations.

- Corrects a potential locally exploitable integer overflow vulnerability in the Network Time Protocol. This potential vulnerability could lead to clients receiving an incorrect date/time offset, resulting in an incorrect date/time on the client.  
SRT4718 - NTP (Severity - High)
- Modifies the ip6\_setup script to limit some Mobile IPv6 questions to LAN interfaces only.
- Corrects a problem in which ftpd core dumps when a 1000 or more directories are present.
- Adds support for ifconfig [inet6] delete [abort] for IPv6 connections.
- Fixes a problem in the implementation of the RIPng protocol that prevents IPv6 routes from being deleted as expected.
- Resolves intermittent core failures in gated.
- Corrects default tape device as /dev/tape/tape0\_d1 for dump and restore as per device naming convention in V5x versions.

## Patch 27021.00

### OSFCMPLRS540

- Corrects default values for YESEXPR and NOEXPR defined in the localedef command and libc to get correct return value from nl\_langinfo(YESEXPR) and nl\_langinfo(NOEXPR)
- Resolves a problem that could cause the rexec() function to hang.
- Fixes the getaddrinfo() routine to work properly when IPv6 is not configured.
- Fixes memory leaks in the libc getipnodebyname routine, which is used by ldapcd. The leaks caused intermittent SSO SIA authentication failures.
- Fixes an issue with the KZPCC backplane RAID adapter device driver (I2O) that causes its logical disk drives to be identified as SCSI devices
- Fixes the mountd daemon to prevent it from becoming unresponsive at a few large sites.
- Fixes a problem in SIA by resetting the mechanism's context pkgind on a successful return of (set|end)\*ent calls.
- Fixes a POSIX standard violation in the strfmon() function. The preceding and following spaces will be padded to return value to make equal length between positive and negative values.
- Fixes a security issue with the C library routine getnameinfo().
- Fixes getnameinfo() to display an IPv4 address instead of an IPv4 mapped IPv6 address when the BIND mapping does not exist.
- Fixes a POSIX standard violation in the wcstod() function. An incorrect pointer was set to the endptr parameter of wcstod() for cases where no conversion was made.
- Fixes the swprintf() function to return the correct value if it detects an invalid wide-character.
- Fixes the ypwhich -m command to prevent RPC timeout error messages.
- Corrects a security issue in which rsh and other rcmds incorrectly report ESUCCESS when the remote side of a connection terminates before fully establishing a connection.
- Allows the auditing of login and su events based in part on the contents of user profiles (for Enhanced Security), the prevailing auditing characteristics of the originating process, and the system-wide audit mask. Previously, only the system audit mask was referenced.
- Fixes a problem with floating point data inconsistencies in threaded applications.
- Corrects RPC-based servers' handling of ill-formed TCP connections.
- Prevents segmentation faults when sia\_ses\_init is passed a malformed argument vector.
- Corrects a potential security vulnerability that may result in a Denial of Service (DoS). This may be in the form of local and remote security domain risks.  
(SSRT2384 rpc — Severity - High)
- Fixes a problem in which the home directory and login shell attributes for a user account are not supplied to the audit daemon for authentication failures.



- Fixes an extended regular expression problem where the interval expression {m,n} is handled incorrectly.
- Fixes a problem from pre-Version 5.0 releases in the way the libc mktime() function handles potentially ambiguous tm struct times that fall within a backward clock shift and that have an initially negative tm\_isdst value.
- Fixes various problems in the libc functions getdate(), strptime(), callrpc(), strncasecmp() and fork().
- Fixes a problem in the libnuma function ncreate() and the system header <sgtty.h>.
- Fixes various problems in the dbx and object file tools dbx, ostrip, strip, mcs, dis, cord, file, and stdump.
- 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.
- Corrects a problem in which some networking applications, especially X.25 and X.29, stopped working as expected because of interactions with security-related fixes and how the fstat() function behaves on their sockets.
- Corrects a potential security vulnerability that may result in nonprivileged users gaining unauthorized access to files or privileged access on the system. This may be in the form of a local and remote security domain risk.
- Fixes a fatal assertion error reported by pixie, hiprof, third spike, cord, uprofile and odump object file tools for some executable files linked at optimization level 2 (-O2) or greater.
- Corrects 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 number of regular expression problems in multibyte locales and a possible hang problem with complex regular expressions.
- Delivers version 3.07.10 of the Tru64 UNIX assembler, which fixes a problem encountered in version 3.07.09, wherein the assembler incorrectly treats octal constant data as if it were decimal.
- 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.
- Corrects problems with name resolution when an error is encountered during the processing of the local host files.
- Fixes a yacc stack overflow error in the Tru64 UNIX assembler.
- Fixes a fatal error in /usr/bin/spike.
- Fixes problems such as segmentation faults caused by the strxfrm() function running on the French locales.
- Changes the use of the configuration file /etc/svc.conf to /etc/nsswitch.conf to allow netgroup data to be provided from LDAP, rather than only from NIS.
- Fixes a performance problem in the libc mktime() routine.
- Corrects a problem in which the rewind() function would fail to reposition to the beginning of a file.
- Corrects a failure in the safe\_open() routine that caused symbolic links given by a relative path from the current working directory sometimes to give ENOENT errors incorrectly.
- Corrects an odd, unexpected error message that may be printed by rsh or rlogin commands.
- Corrects a problem where the telnet command causes unnecessary delays when an IP address is supplied as a command-line argument.

- Corrects a problem in which a DNS resolver routines never time out if interrupted by signals.
- Provides an RFC3542 compliant implementation of IPv6 Advanced API.
- Fixes a nonconcurrency issue for multithreaded applications calling popen() and certain "FILE \*" routines such as fread().
- Fixes a deadlock condition in multithreaded applications that call fork() and other libc callback routines such as exit handlers, `__fini_*` routines.
- Corrects a condition in which multiple "Sorry" messages are issued by the su command when multiple SIA mechanisms are in use (as when LDAP is configured for user accounts).
- Improves the performance of times(3).
- Enhances the fuser command to provide a cluster-wide query capability. A revised fuser(8) reference page describes this enhancement.
- Fixes a problem with glob() returning an incorrect match when directory permissions issues exist.

#### Patch 27022.00

##### *OSFDCMT540*

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

#### Patch 27023.00

##### *OSFDCMTEXT540*

- 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.

#### Patch 27026.00

##### *OSFDOSTOOLS540*

- 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 several security vulnerabilities.
- Changes mcopy and mwrite so they can overwrite existing files.
- Changes mtools so they can print appropriate error messages for nonprivileged users.
- Standardizes the mformat prompt.

#### Patch 27028.00

##### *OSFENVMON540*

- Fixes the envmond daemon to use EVM events for monitoring environmental events.
- Fixes envmond to report redundant power supply restoration on systems that support environmental monitoring through the sysconfig command.
- Modifies the environmental monitoring facilities /usr/sbin/envmond and /usr/sbin/envconfig to support the AlphaServer GS1280 system.
- Updates envmond to ensure that the correct EVM events are being sent at the correct time.
- Corrects a problem in which the environmental monitoring daemon does not display the debug message "proceed ..." when the system is booted.
- Corrects a problem in which envmond sets rc.config variable ENVMON\_HIGH\_THRESH to a value read from kernel module if ENVMON\_HIGH\_THRESH is null.
- Fixes issues with temperature threshold handling in envmond.

- Corrects a performance problem on systems with many sensors by reducing the polling frequency.
- Prevents a long boot pause on AlphaServer GS1280 systems with large disk and sensor counts.
- Fixes issues with envmond with respect to log messages and the execution of shutdown scripts.

#### Patch 27030.00

##### *OSFEXAMPLES540*

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

#### Patch 27031.00

##### *OSFEXER540*

- Corrects the memory exerciser user.syslog message from "Started" to "Stopped".
- Corrects a message from memx.

#### Patch 27034.00

##### *OSFHWDATABASE540*

- Enhances hwmgr show functionality for scsi path information.
- Provides for FCP-2 Link Level Error Recovery for tape operations so that tape I/O can be successfully performed during periods of link errors, noisy lines, bit errors, and lost or corrupted data frames.
- Enables Tru64 UNIX to work with tape devices that do not support non-tagged commands.
- Fix for netstat trying to resolve the IP address 0.0.0.0 (INADDR\_ANY) to a hostname.
- Fix to netstat, kdbx displaying partial kernel addresses.
- Fixes several I/O error handling and error reporting problems in the Tru64 Emulex Fibre Channel driver that can result in inefficient error handling or misleading error log entries.
- Updates the bcm driver to V1.0.22 to fix issues with IPv6 and SNMP
- Fixes many small problems with the dsfmgr command.
- Fixes the hwmgr command to show path state correctly.
- Fixes numerous issues in the driver for DEGXA Gigabit Ethernet adapters, including the DS25 onboard 10/100/1000 port.
- Corrects a problem in which information from the hwmgr -view transaction -cluster command for a node on a cluster may not be displayed.
- Corrects some command-parsing irregularities in hwmgr that may cause options like -category and -cluster to be confused.
- Fixes a problem in which the display for the hwmgr -show name command is not aligned properly for the name field.
- Adds IEEE 802.1Q Virtual Local Area Network (VLAN) support for the following:
  - DEGPA
  - DEGXA
  - DE50x, lan\_common.h
  - DE60x
- Fixes a problem in the alt driver for DEGPA Gigabit Ethernet adapters. This problem affects all Tru64 UNIX systems containing DEGPA network interfaces.
- Fixes a problem with scu where a mismatch between expected and found data causes incorrect data to be displayed.
- Provides additional support for Ultrium 2 SCSI tape drive.
- Ensures proper compilation of the DDR database.

- Modifies netstat and ifconfig so that when a MAC address is printed, it is printed using 2-digit hex octets with leading zeros.
- Fixes a problem in which Smart Array 5300 logical volumes are counted as RAID controllers.
- Adds dumprmt.msg for remote dump/restore messages. This new message catalog file is used in both rdump and rrestore programs.
- Corrects a potential system crash when shutting down after using a DAPBA or DAPCA ATM adapter.
- Fixes a problem with non-U.S. USB keyboards used in non-U.S. locales in which the keyboards are treated as U.S. keyboards by the operating system.
- Fixes various problems in the bcm driver for DEGXA Gigabit Ethernet which can cause crashes.
- Fixes the dump command to recognize LSM volumes correctly and to not report random information when an error has occurred.
- Fixes a condition that may cause a panic if the Xserver is stopped.
- Changes the fwupgrade command to allow the specified firmware update image to be located on a BOOTP server in a connected network.
- Fixes several problems in the bcm driver for DEGXA Gigabit Ethernet adapters, including the following:
  - A condition that causes the driver to incorrectly report data overruns
  - A condition that prevents DEGXA2-SA modules from being recognized
- Fixes two problems in the tu driver for DE5xx 10/100 MB Ethernet adapters.
- Adds code to print greater than 61 UNIX domain sockets.
- Changes file read errors from /dev/kmem to ignore and continue in a running system.
- Fixes a problem in the alt driver that prevents DEGPA from being used with DE50x or DE60x adapters in a LAG set.
- Increases the default limit of DLI packets to 16 K and makes the limit tunable. This corrects a problem in which attempts to sent packets larger than 5000 bytes (jumbo packets) can fail.
- Updates DDR for the MSA array controller and for other future devices.
- Adds recognition for possible future devices.
- Fixes various problems in the bcm driver for DEGXA Gigabit Ethernet adapters.
- Updates ICMP code to add redirect timeout support and non-modifiable route table entries.
- Provides modifications for emx hardware.
- Provides device support for the SDLT160/320 tape drive.
- Fixes various problems in the ee driver for DE60x.
- Sets a software limit to prevent serious performance problems that can occur when TCP connections that have an rate-limit enforced by a downstream network device overrun the device.
- Corrects a problem in which /sbin/ddr\_config does not accept values for ReadyTimeSeconds larger than 255. The new limit is 86400 seconds (24 hours).
- Fixes problems with NUMA disk statistics
- Fixes a KMF problem that can happen if some of the nodes in cluster are rebooted and a device is shared by all the nodes.
- Changes the CAM subsystem message that is printed to the error log on a recovered read error from "bad block number" to "block number."
- Corrects a problem that can cause a core dump when the fwupgrade command is executed with qualifiers.
- Enhances HP-XP Array controller support.
- Corrects several potential security vulnerabilities where, under certain circumstances, system integrity may be compromised. These may be in the form of improper file or privilege management.

- Fixes an I/O hang condition on Fibre Channel.
  - 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.
  - Updates the Radeon graphics driver to not reject multi-headed configurations.
  - Fixes a memory fault condition in the emx driver that occurs when responding to an inquiry command from a remote port in the fabric.
  - Updates DDR for the SDLT600, DLT VS, SuperDLT1, SDLT320, and VS series of tape drives.
  - Modifies a script that sets up entries in /dev to solve the following problem: When creating a cluster, certain old style /dev entries should be removed, but in some configurations they are not.
  - Provides support for the Ultrium 2 SCSI tape drive.
  - Corrects a condition that causes bootstrap address collisions after console heap\_expand uses more memory and extends into address range used by the kernel.
  - Updates the tu driver to support Ethernet Multicast addresses larger than 512, which is necessary to support a large number of IPv6 addresses
  - Updates the alt driver to V2.0.20 to fix issues with IPv6, NFS performance and SNMP
  - Fixes an improper handling of domain, area, and fabric RSCNs, or Registered State Change Notifications by the emx driver that results in the nondetection of path failures to storage devices in the fabric.
  - Fixes and improves the mcutil program by correcting how bus resets are handled by the program and enhancing its error reporting capabilities.
  - Fixes a problem where Ultrium 2 tape drives do not use hardware compression.
  - Updates the ee driver to V1.0.27 to fix issues with IPv6 and SNMP
  - Adds path failure detection to the Tru64 Emulex Fibre Channel driver for situations involving hung N\_Ports of a storage controller in a switched fabric environment.
  - Updates Mobile IPv6 code to be compliant with the latest RFC.
  - Modifies the DDR database to add support for the HP DAT72X6 changer.
  - Updates the bcm driver to V1.0.23 to fix issues with Jumbo frames.
  - Corrects a problem with HSZ events that cause I/O errors rather than doing the appropriate retries so the I/O can be successful.
  - Adds read/write ability to XP control ports.
  - Corrects an issue in which tape devices that have no ddr entry would not get compression correctly enabled or disabled as appropriate to the device special file that was used.
  - Fixes a problem that prevents keyboard input during interactive booting (SRM boot flag i) when in graphics mode on GS1280, ES80, and ES47 AlphaServers.
  - Fixes mass storage disk subsystem handling of specific, important events reported by the HSG80 and HSG60 array controllers.
  - Corrects default tape devices as /dev/tape/tape0\_d1 for dump and restore as per device naming convention in V5x versions.
- Changes I/O retry processing when a command time out is encountered. With this change, the retry code attempts to take advantage of other active paths during retries if command time outs are encountered.
- Corrects a problem with the DAT160 tape drive in which tape devices with no ddr entry would not get compression correctly enabled or disabled as appropriate to the device special file that was used.
  - Adds ddr.dbase support for the DAT160 tape drive.
  - Corrects A potential security vulnerability.

## Patch 27035.00

### OSFHWBIN540

- Prevents the state machine performing Bad Block Replacement on a disk for which BBR is disabled.
- Fixes a rare panic during boot on GS1280/ES80/ES47.
- Fixes a "lock\_fault" panic that can happen during system startup and shutdown.
- Corrects the handling of certain disk "not ready" conditions to prevent long I/O stall times.
- Provides performance enhancements based on the vm\_overflow feature added in V5.1B-3.
- Changes the way page migrations occur on a NUMA system to address poor performance due to excessive paging.
- Corrects a problem with pagetable page allocations that can leave a thread waiting indefinitely during a fork operation.
- Corrects the cause of "ubc\_wire: hash failed" panics on non-NUMA systems.
- Correct the cause of "not wired" panics with System V shared memory and bigpages.
- Fixes an issue in the VM subsystem, wherein a page that is not managed by VM is incorrectly identified as being managed by VM.
- Provides enhanced hwmgr show functionality for scsi path information.
- Fixes a problem in which crash dumps to Fibre Channel swap devices do not always succeed.
- Allows for control ports to be deleted by the hwmgr utility.
- Fixes a cluster deadlock/hang issue that occurs when a new device is discovered.
- Fixes a problem in which the "hwmgr refresh scsi -all" command does not always remove stale paths.
- Corrects the cause of a process hang and system panic that occurs when using System V shared memory and asynchronous I/O.
- Provides preferred path support for active-active, asymmetric storage devices.
- Improves tape read/write command reliability in SANs by allowing the use of FCP-2 Link Level Error Recovery (LLER) for read and write tape commands if the destination Fibre Channel port supports LLER.
- Enables Tru64 UNIX to work with tape devices that do not support non-tagged commands.
- Addresses a potential I/O performance bottleneck in which the tape driver may select the same path for all tape drives in a system containing multiple HBAs and multiple tapes. With this change, the tape driver can assign different paths to different tape drives to improve tape I/O performance.
- Addresses a NUMA memory initialization issue that can lead to system instability.
- Allows systems with multiple paths to a large number of devices to boot faster.
- Fixes the cause of the "pmap\_pagemove" panic that occurs with some third party drivers.
- Fixes a problem in which the changer code causes a kernel memory fault.
- Fixes an issue with the KZPCC backplane RAID adapter device driver (I2O) that causes its logical disk drives to be identified as SCSI devices.
- Makes adapter error messages more descriptive.
- Corrects a memory leak.
- Fixes the problem with the hwmgr utility displaying the wrong output for network device MTU.
- Addresses an issue in which the cam changer driver does not unwire user pages when a path fails.
- Fixes a problem of system hangs that occur due to combined memory leak and DMA resource leak in the EMX driver during the processing of an UNTAGGED request like device/lun reset.

- Fixes a problem with KZPAC (SWXCR controller) that occurs when an open() issued on a deleted unit succeeds and a read I/O request results in a system crash.
- Allows systems with multiple paths to a large number of devices to boot faster.
- Changes the way the aha\_chim OSM driver handles timeouts. Previously, timeouts were handled by registering a timer in the kernel callout queue for each I/O that timed out. As a result, a burst of aha\_chim timeouts could flood the kernel callout queue. The new method times the I/O itself, so no entries are put in the kernel callout queue.
- Ensures that the alt firmware is stopped before freeing buffers when the alt card is re-initializing.
- Fixes a problem with the alt driver regarding receiving 9000-byte packets.
- Provides transmit/receive flow control settings and driver revision to the bcm Gigabit Ethernet driver boot time startup messages.
- Provides a fix for the proper operation of DS25 onboard bcm LEDs.
- Fixes an issue in the error recovery procedure of CISS adapters.
- Improves failure detection handling due to command timeouts.
- Fixes a potential race condition that causes a kernel memory fault by the changer code.
- Allow systems with multiple paths to a large number of devices to boot faster.
- Fixes inappropriate handling of failures during tape changer opens.
- Corrects a problem in which the consvar command can panic or fail without error when setting some console environment variables.
- Fixes problems in the aha\_chim (KZPEA) driver.
- Corrects a potential hang issue with tapes when their SCSI address changes.
- Enables FCP-2 Link Level Error Recovery for tape operations.
- Corrects the handling of LS\_RJT messages in the emx driver.
- Corrects the cause of a kernel memory fault panic in the KZPEA driver.
- Fixes a simple lock fault in wakeup\_async\_waiter.
- Fixes a problem with the environmental monitoring daemon on partitioned AlphaServer GS80/GS160/GS320 systems.
- Provides new EMX driver hardware attributes.
- Improves selection timeout handling in RAID services.
- Fixes a simple lock fault in cdisk\_bbr\_comp.
- Fixes an issue related to aha\_chim driver blocking in interrupt context.
- Fixes the cause of a kernel memory fault that occurs in the ctape\_generic\_pass\_thru routine when bus resets were detected.
- Prevents system hangs due to certain internal Fibre Channel adapter failures.
- Fixes an issue in which the hwmgr utility displays an incorrect MTU size for a network interface.
- Fixes several I/O error handling and error reporting problems in the Tru64 Emulex Fibre Channel driver that can result in inefficient error handling or misleading error log entries.
- Causes an error to be logged to binary.errlog when a LUN contraction is detected
- Causes "Tru64 AHA\_CHIM" to be returned for the SIM vendor ID.
- Allows the hwmgr redirect scsi command to work with lockmode 4.
- Fixes device recovery timing by improving the chances for a problem disk to recover.
- Allows the first path to be removed with hwmgr if it is stale.
- Provides environmental sensor support for Superdome power supply sensors.
- Provides support for the CPU offline capability on AlphaServer GS1280 systems, as required for Capacity on Demand.
- Updates the bcm driver to V1.0.22 to fix issues with IPv6 and SNMP.

- Enables SmartArray 5300 controller hardware events to be logged to the `binary.errlog` during boot time. This is useful in diagnosing logical volume state change and physical drive hotswaps that can occur while the system is not booted.
- Fixes a simple lock panic in the floppy diskette driver.
- Prevents the memory controller from starting on titan and tsunami platforms with aluminum ev68 CPUs.
- Fixes a problem with the Smart Array driver that could cause a system hang to occur during error recovery when I/O is active.
- Fixes a `consvar -s bootdef_dev` failure with KZPCC.
- Corrects a problem in which systems configured with VX1 graphics card do not return to console when the halt button is pressed, thereby making the console unusable.
- Adds an event to indicate that the soft or hard error count has changed on the device identified in the event.
- Corrects a potential deadlock in the hardware configuration subsystem.
- Fixes a problem where, when using `hwmgr` to delete a component, a "DELETE\_COMMIT: Cannot fetch name." message may be displayed on the console. This problem can be seen frequently in a cluster environment when the component being deleted does not exist on the system.
- Fixes numerous issues in the driver for DEGXA Gigabit Ethernet adapters, including the DS25 onboard 10/100/1000 port.
- Corrects a problem in which incorrect values for `LONG_MAX` and `LONG_MIN` were displayed when using the hardware manager to show attributes.
- Introduces type checking of attributes when registering components with the hardware manager.
- Corrects a problem where after entering a `hwmgr -redirect SCSI` command and rebooting, the system only boots to single user mode with the following error displayed:
 

```
bcheckrc: Device Naming failed boot configure
or verify Please correct the problem and continue or reboot INIT:
SINGLE-USER MODE
```
- Corrects invalid `hwmgr show component` inconsistency.
- Addresses a problem encountered when mounting cluster root if the cluster root domain devices are private to different cluster members. Currently, you cannot boot your cluster. It will hang. With this fix, your cluster will boot with a warning to the console. This configuration is not recommended; however the cluster should not be unbootable. Currently, this is with respect to non-LSM cluster root domains.
- Prevents the hardware management cluster database from being reset.
- Adds IEEE 802.1Q Virtual Local Area Network (VLAN) support for the following:
  - DEGPA
  - DEGXA
  - DE50x, `lan_common.h`
  - DE60x
- Removes a restriction in which dynamic VMEbus device drivers can probe only one controller per driver.
- Corrects a potential floating point register inconsistency.
- Corrects 3D client hangs when using the Radeon graphics card.
- Corrects a problem in which a kernel memory fault sometimes occurs if a USB keyboard or mouse does not respond quickly enough. This KMF can occur during boot or soon after a USB keyboard or mouse is connected. Any device can trigger this, though it is neither predictable nor common.
- Fixes a problem in the alt driver for DEGPA Gigabit Ethernet adapters. This problem affects all Tru64 UNIX systems containing DEGPA network interfaces.



- Fixes a condition that causes a system hang when using open3D over the AGP bus on an AlphaServer GS1280.
- Fixes a problem that can cause an ES45 to hang if the Xserver is restarted or the system is rebooted without a power cycle when using the Radeon AGP graphics device.
- Fixes a problem with USB hubs (or any other bus device) that occurs when they are removed from a running system.
- Fixes a performance problem that occurs when doing wiring on gh\_chunks memory.
- Fixes multiple problems affecting a system with peripheral USB hubs attached, as well as problems that might occur when moving or adding USB host adapters.
- Fixes a situation in which mounting a valid CD-ROM for the first time fails with the message "No valid file system exists on this partition," although subsequent mounts of the same CD-ROM work as expected.
- Fixes a kernel memory fault for systems that contain more than 8 IDE/ATA buses.
- Corrects rounding errors for vm attribute vm\_bigpg\_thresh.
- Corrects the handling of bad pages when big pages are enabled.
- Fixes the cause of "page mapped" panics when using mmap calls with dev/mem to access free big pages.
- 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.
- Adds support in the platform code to handle MSI capable adapters. AlphaServer GS1280 systems support option cards that require MSI capabilities.
- Installs the V1.07 release of the ciss driver, which is the minimum version required to support the Smart Array 5300 controller.
- Corrects a problem in which a path event can cause hang in cdisk\_online during disk open of HSG80.
- Provides additional environmental support for the DS20L platform.
- Allows multiple VX1 graphic cards to be configured in a separate I/O box system.
- Adds support to get live status information for air movers and power supplies on AlphaServer GS1280 systems and to log intrusion packets to the error log.
- Adds support for CPU indictment on AlphaServer GS1280 systems.
- Corrects problems with the time of year (TOY) clock.
- Fixes an IDE/ATA bus hang caused by attempting to complete raw odd byte DMA transfers to or from IDE/ATAPI devices.
- Addresses small memory leaks within the kernel that occur infrequently.
- Adds sysconfig tunable attributes for AlphaServer ES45 environmental monitoring.
- Addresses an issue in which a NULL Inquiry data causes a "Device has no 'name' " error and possible I/O stalls.
- Corrects the cause of a delete\_pv\_entry panic when kernel virtual address space has high usage.
- Fixes a problem in which Smart Array 5300 logical volumes are counted as RAID controllers.
- Address an issue in which AdvFS domain panics occur during HSZ and HSG failovers.
- Improves I/O performance by reducing kernel locking overhead.
- Fixes a small memory leak in Power Management code.
- Prevents unnecessary retries on an HSG80 when fail unit attention with ascq = 0xf002 and returns proper error to higher layer.
- 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.
- Corrects a potential system crash when shutting down after using a DAPBA or DAPCA ATM adapter.

- Prevents an IDE bus hang caused when issuing a play audio track command from scu to an ATAPI CD-ROM containing an enhanced CD.
- Installs Version 1.08 of the ciss driver.
- Fixes a problem with non-U.S. USB keyboards used in non-U.S. locales in which the keyboards are treated as U.S. keyboards by the operating system.
- Fixes various problems in the bcm driver for DEGXA Gigabit Ethernet that can cause crashes.
- Corrects a problem in the marvel\_pfm driver where xmesh and bmesh incorrectly reported 100 percent utilization for IO7 ports held in reset. Previously, this had to be corrected by using the mvfi test program.
- Makes it possible for CPUs that do not take interrupts (either directly from the attached IO7 or indirectly from IO7s attached to offlined CPU) to always be allowed to be offlined.
- 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 recursive panic situation on the ES47 platform when a double bit memory error is detected.
- Corrects a condition to prevents erroneous "ccfg\_MakeDeviceIdentWWID: Invalid device ID" messages from being generated.
- Fixes a problem that may cause a panic if the Xserver is stopped.
- Fixes a process hang in ubc\_common\_lookup.
- Corrects a problem in which the BBR code logs all error messages as soft error, even if the error was not recovered and it failed to do the bad block replacement.
- Fixes a condition that causes a rare hang in the hardware configuration subsystem.
- Resolves a problem in which some DE50x network interface cards, under specific circumstances, may not send gratuitous arp packets.
- Fixes the re\_ioctl() cases DIODCMD and DIODCDB where cmd transfer size has been changed to avoid kernel memory fault.
- Changes the fwupgrade command to allow the specified firmware update image to be located on a BOOTP server in a connected network.
- Suppresses an erroneous console warning message that may be provided when cluster root is under LSM control. The warning "WARNING: cluster root devices are on private buses!" may be erroneously output when cluster root is under LSM control. LSM does not support such configurations.
- Fixes a problem in which the CAM I/O subsystem can cause a kernel memory fault or system hang when the subsystem is low on memory.
- Installs Version 1.09 of the ciss driver.
- Installs Version 1.10 of the ciss driver.
- Modifies console callback code to allow users to use upper and lower case variable names for known console environment variables. This patch is required for update installations on EV7 based platforms.
- Corrects the creation of console boot device strings for devices on subordinate buses.
- Fixes the garbage character that sometimes appears when requesting the name of a boot device via consvar.
- Fixes several problems in the bcm driver for DEGXA Gigabit Ethernet adapters, including the following:
  - A condition that causes the driver to incorrectly report data overruns.
  - A condition that prevents DEGXA2-SA modules from being recognized.
- Fixes two problems in the tu driver for DE5xx 10/100 MB Ethernet adapters.
- Clarifies the "LIDs do not match" error message by displaying the values that do not match.
- Corrects a race condition that may result in hung disks under certain circumstances, for example, after a SCSI reset.

- Corrects problems in which tape devices become unavailable, do not respond, report unrecoverable errors, or cause a kernel memory fault.
- Corrects problems in the aha\_chim driver that could result in bus hangs, panics and inappropriate access of freed memory during high rate of bus resets.
- Modifies the changer device driver so that it reports the manufacturer ID attribute as could be seen with the hwmgr -view devices command.
- Fixes a problem in the alt driver that prevents DEGPA from being used with DE50x or DE60x adapters in a LAG set.
- Fixes the cause of the panic "dg: unwiring."
- Fixes a problem in the KZPCA itpsa driver that occurs when a SCSI target presents multiple LUNs.
- Prevents problems in the USB subsystem, including memory leaks, data inconsistencies, and USB device configuration problems.
- Installs Version 1.11 of the ciss driver.
- Fixes problems in the USB driver layer, including minor performance degradations and device failures.
- Reduces the delay with hwmgr of EV7 based machines by reducing the number of calls to the console to update sensor data.
- Fixes several IPMI-related problems, including the following:
  - Erroneous fields in 686 OS-detected environmental machine check logout frame.
  - Unusually large number of 686 sensor timeouts with heavy system load.
  - IPMI always reporting -48v sensors as broken, seen as "redundant power supply failed" messages
  - An IPMI memory leak
- Restores the well known panic strings "Processor Machine Check" and "System Uncorrectable Machine Check," which had been replaced by more specific error strings. With this patch, both the old panic strings and the more specific ones will be issued.
- Adds recognition for possible future devices.
- Adds the capability for KZPCA devices to work with SCSI devices that only support asynchronous data transfers.
- Fixes a problem in which the SDLT media causes bus resets.
- Corrects problems that can result in tape devices not responding or in a kernel memory fault when the FNDELAY flag is set in tape open and tape ioctl.
- Addresses an issue in which I/O may not complete under certain circumstances.
- Installs version 1.12 of the ciss driver to correct a condition that causes an AlphaServer GS80 to hang following an Smart Array 5300 reset.
- Corrects lockmode 4 problem in cdisk\_event\_notify.
- Corrects a condition in which 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.
- Corrects an illegal instruction that causes a rare system crash.
- Fixes a problem in which the user cannot connect to a printer on the parallel port after cancelling a print job.
- Adds support for EV7 class machines (for example, the AlphaServer GS1280) to allow the use of processors having different speeds in the same chassis.
- Corrects a condition in which the KZPEA firmware fails to correctly handle file marks with odd byte transfers.
- Fixes a problem in the PCI bus code that can prevent some functions on a multifunction PCI card from being configured.

- Fixes a problem with the USB keyboard driver that affects keyboard operation when the X server is not running. Specifically, the autorepeat function it may appear jumpy, or keys may appear to press themselves after they are released.
- Fixes a problem relating to gh\_chunks allocation on some configurations.
- Fixes various problems in the bcm driver for DEGXA Gigabit Ethernet adapters.
- Provides the logging of an informational event to the errorlog when the peripheral driver is unable to get a user specified IDENTIFIER for a storage array device.
- Fixes CAM errors that occur when opening CD devices containing blank media.
- Corrects a problem in which a read or write operation to a changer device creates an unkillable process.
- Fixes a problem that can cause a kernel memory fault during boot if a Fibre Channel LSM boot disk is not discovered during the device probe.
- Suppresses an in\_cksum() console debug message.
- Fixes a memory fault condition in the user agent driver that occurs when multiple threads issue I/Os on different CPUs.
- Fixes jitter problems in 24-bit depth on VX1 graphics cards with certain date codes.
- Fixes the cause of a system crash that can occur if the hwmgr command deletes a disk while it is in recovery.
- Increases BOOTP error timeout values for RIS kernel installations.
- Updates the values in the header of the PCI subpackets.
- 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.
- Allows multiple PCI-X devices on a system to use Message Signaled Interrupts (MSI).
- Reduces the existing limitation on the number of CPUs that can be taken off line by allowing the other CPUs in the system to handle the interrupts coming from up to two IO7s.
- Corrects a "simple\_lock: time limit exceeded" issue that can occur in the CAM I/O subsystem when MCS locking is disabled.
- Fixes a problem that turns off the reporting of correctable errors forever on any CPU, except CPU 0, once the throttling of correctable errors has begun.
- Causes the posting of environmental events in the event of a machine check on EV7-based platforms.
- Corrects a problem on AlphaServer ES45 platforms that resets sysconfigtab settings for the attribute titan\_sys\_ps\_hotswap after a system reboot.
- Provides modifications for possible future emx hardware.
- Provides support for the Philips USB controller, which is shipped on some AlphaServer GS1280 systems. Without this patch, the Philips USB controller may fail to detect and configure devices below it.
- Fixes various problems in the ee driver for DE60x Ethernet adapters.
- Corrects a problem in which /sbin/ddr\_config does not accept values for ReadyTimeSeconds larger than 255. The new limit is 86400 seconds (24 hours).
- Fixes problems with NUMA disk statistics.
- Fixes a KMF problem can occur if some nodes in cluster are rebooted and a device is shared by all the nodes.
- Changes the CAM subsystem message that is printed to the error log on a recovered read error from "bad block number" to "block number."
- Provides hardware support for the AlphaServer DS15/TS15 platform.
- Corrects a race condition that can cause a process to hang during disk error recovery failure processing.
- Fixes a hole in disk I/O handling that can cause a panic in rare circumstances.

- Corrects a problem that causes a system panic while running applications that open a RAID device and the faulting routine is control\_port\_open.
- Fixes an I/O hang condition on Fibre Channel.
- Fixes a rare case in which the target of a ladefunc-invoked routine (via the call func command) aborts with a segmentation fault while dereferencing the gp register.
- Corrects a problem in which unconfigured PCI devices can cause a panic during boot.
- 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.
- Corrects an address problem with hwmgr delete while a SCSI scan is in progress.
- Updates the Radeon driver to not reject multiheaded configurations.
- Removes an unwanted test text message.
- Fixes a memory fault condition in the emx driver that occurs when responding to an inquiry command from a remote port in the fabric.
- Corrects a problem that reduces performance or causes a device to switch controllers unnecessarily when active paths are improperly treated as standby or vice versa.
- Fixes a problem in which locks used for disk I/O are not released during rare error conditions.
- Fixes a problem with disk I/O barrier code that was not executed when it was needed.
- Corrects a condition in which a panic can occur when unrecognized PCI devices are probed with a PCI class code of 255. The problem is probably limited to very old or specialized devices.
- Corrects a HSZ70 controller failover failure that generates the message "Logical unit not ready, cause not reportable."
- Changes the way reservation-conflict errors are handled in a cluster.
- Prevents a potential panic that generates the message "memory\_test=partial" or "memory\_test=none."
- Corrects a problem in which certain older tape drives (such as the TZK50 and TSV07) would produce only errors when they were accessed.
- Causes the AlphaServer GS1280 to accurately report and handle fatal memory errors if the associated CPU is off line.
- Extends the hardware managers for CPU entries to include the attributes on a GS1280 platform to include the CPU chip revision number and the cache size.
- Fixes a condition that causes hung I/Os in systems with multiple RADs.
- Updates the tu driver to support Ethernet Multicast addresses larger than 512, which is necessary to support a large number of IPv6 addresses
- Corrects a problem that causes the command sysconfig -r hwc hwc\_print=5 to generate a kernel memory fault.
- Updates the alt driver to V2.0.20 to fix issues with IPv6, NFS performance, and SNMP.
- Fixes an improper handling of domain, area, and fabric RSCNs (Registered State Change Notifications) by the emx driver that results in the nondetection of path failures to storage devices in the fabric."
- Fixes the reporting of device monitoring events and hardware errors during disk recovery from the disk driver to the binary errlog.
- Corrects a problem that occurs with a device or bus reset during the execution of a command to a media changer device, like a tape library, and causes the system to panic with the message "PWS\_CCB\_QUE\_REMOVE: ccb not on any list."
- Prevents a cluster panic when installing new RAID hardware before OS support is installed on the cluster.
- Fixes a problem in which an I/O operation is prematurely treated as complete, which can cause a panic in XPT or a storage driver.

- Improves error handling and I/O timing.
- Fixes hangs that can occur prior to disk spinups.
- Improves system recovery when media errors occur.
- Eliminates a condition that causes a panic to occur when removing sensors on a ES47 system.
- Fixes a problem with HSG80 controllers due to bad count for persistent reservation keys.
- Improves the performance of systems that perform heavy file I/O.
- Changes the way that memory is marked as bad on the AlphaServer GS1280 in order to prevent a recursive crash on reboots. (A crash is expected as a result of a double-bit-memory error machine check.)
- Fixes the errant binary .errlog entry "Status = CMP but resid not NULL Possible Software Problem - Impossible Cond Detected" from the peripheral disk, tape, and changer drivers.
- Updates the ee driver to V1.0.27 to fix issues with IPv6 and SNMP
- Corrects a potential race condition in the cam\_tape driver.
- Fixes a problem that causes a system to hang when swapping.
- Corrects the cause of a simple lock timeout panic during I/O barriers.
- Adds path failure detection to the Emulex Fibre Channel driver for situations involving hung N\_Ports of a storage controller in a switched fabric environment.
- Corrects the cause of a system crash that occurs when deleting disks using hwmgr.
- Fixes the cause of certain panics by increasing the timeout for CPU onlines.
- Provides better error handling for the MSA1000.
- Fixes a problem in which a node does not properly find the server for a disk but incorrectly fails an open with ENODEV.
- Installs Version 1.13 of the ciss driver.
- Provides changes that allow for the proper usage of Smart Array controllers that are behind PCI bridges.
- Provides support for new PCI svid for the ATM controllers 155 mmf, 155 utp, and 622 mmf.
- Fixes the 'hwmgr edit scsi command.
- Fixes the cause of a "Simple lock owned" panic seen on a KZPEA during a cluster boot.
- Updates the bcm driver to V1.0.23 to fix issues with Jumbo frames.
- Corrects a condition in which the firmware FRU table and Configuration Tree are not logged to binary errorlog during boot on AlphaServer ES47/ES80/GS1280 systems with less than 4 GB of memory.
- Fixes the cause of a system crash that occurs when running with lockmode=4.
- Corrects a problem in which the hwmgr command's display of DS15 environmental sensors may indicate that they are in the warning state when they are actually in the fault state.
- Corrects a condition that prevents setting a high temperature threshold (warn and fault) for hwmgr and envmond on DS15.
- Fixes a potential disk-identification delay when doing disk cloning.
- Corrects a problem in which the hwmgr -delete command can cause a panic when directed at a sensor.
- Adds a ddr flag to prevent failing I/O when ASCQ\_LUN\_NRDY\_MAN are received from HSZ devices.
- Adds read/write ability to XP control ports.
- Corrects a problem with a sensor error that can indicate a false over-temperature condition on DS10/DS10L and TS10 systems.
- Fixes a kernel memory fault panic problem with the KZPCC RAID adapter.
- Fixes a panic that can occur during asynchronous event notification by a mass storage driver.

- Corrects a problem in which AlphaServer GS1280/ES80/ES47 systems may report false over-temperature conditions if the temperature drops to (or below) 9 degrees Celsius (48 F).
- Allows hwmgr -redirect to work with lockmode 4.
- Fixes invalid queue entries and system panics following Retries Exhausted.
- Fixes a persistent reduction in SCSI queue depth after a queue full occurs.
- Fixes a race condition in xpt\_async\_event\_thread.
- Fixes an "mcs\_lock:lock already owned by cpu" panic.
- Fixes an issue with the storage subsystem I/O Barrier error handling.
- Fixes USB problems that prevent some USB keyboards and mice from working.
- Fixes POx Uncorrected and Unrecoverable Error panics on EV7 based systems.
- Corrects tag and non-tag I/O handling when the system is booting.
- Corrects a problem in which a multi-threaded application accessing a single tape drive via a single FILE object shared amongst the threads causes a kernel memory fault.
- Fixes a problem with VMAC not working on the bcm driver (DEGXA) without explicitly enabling promiscuous mode.
- Fixes a rare problem in which a "queue full" response during device discovery results in a kernel memory fault.
- Adds additional fast fail functionality to the disk class driver.
- Clarifies the binary error-log entry for medium or hardware errors.
- Correct error logging problems in the emx driver.
- Fixes a problem that causes an "invalid link address" kernel panic in the Alteon gigabit Ethernet driver.
- Fixes the incorrect locking order in alt driver.
- Corrects a problem in which FibreChannel dropped data is not detected during I/O error recovery and so is not retried, thereby causing unexpected and unnecessary errors.
- Corrects a problem in which tape devices with no ddr entry would incorrectly enable or disable compression to the device special file that was used.
- Fixes a problem in which the operating system receives an environmental machine check packet from the firmware, but fails to correctly recognize the sensor that is identified as faulty by the machine check.
- Corrects a problem in which multiple access to changer devices (for example, via Legato/Networker) could lead to a kernel memory fault in the changer driver.
- Fixes problems with DS15 environmental monitoring.
- Fixes mass storage disk subsystem handling of specific important events reported by the HSG80 and HSG60 array controllers.
- Corrects a problem in which multiple processes accessing a changer device can hang in a deadlock situation accessing the changer.
- Improves logging of sense data during error recovery.
- Fixes a problem with stalled or hung I/O's on clusters.
- Corrects the panic, "cmn\_err: CE\_PANIC: elan0: heartbeat pacemaker - flatlined ...." that occurs on AlphaServer SC systems.
- Fixes a panic that can occur when the ee driver momentarily drops its lock.
- Resolves a logic error in hardware-specific platform code for certain boundary conditions.
- Enables IPv6 MTU to be set up to 9000 bytes on gigabit Ethernet.
- Fixes a rare kernel memory fault that occurs when booting an AlphaServer ES80.
- Fixes a timer queue overflow panic with ccmn\_get\_dbuf.
- Fixes a kernel memory fault in hwc\_lookup\_devt\_safe.
- Enables the right mouse button for some two-button USB mice.

- Fixes a kernel memory fault that can occur during boot on some large AlphaServer GS1280 systems.
- Enhances CAM disk driver resiliency.
- Fixes a potential I/O hang with the KZPCC backplane RAID adapter.
- Fixes a rare occurrence where on a shared parallel SCSI bus disks may become inaccessible to hosts under certain fault conditions.
- Fixes the panic "no locks owned by cpu" caused by the bcm driver (DEGXA) when the device is shut down.
- Fixes a problem that causes incorrect parent bus information from being printed for Smart Array 53xx adapters.
- Corrects a lock initialization during boot.
- Prevents a kernel memory fault within control\_port\_open when there are no control ports within the system configuration.
- Fixes a problem and prevents a uagt thread from causing a "simple lock owned" panic in low memory situations.
- Fixes a kernel memory fault in hwc\_lookup\_devt\_safe.
- Improves cluster I/O resiliency during certain error cases.
- Provides enhancements for future versions of XP storage tools.
- Fixes a race condition in the EMX driver that can cause the driver state flags to go bad.
- Fixes the cause of a KMF panic that occurs when closing a removable device with no paths.
- Provides better error handling for new tape devices.
- Fixes an issue in the error recovery procedure of CISS adapters.
- Corrects a problem in which changer applications hang when opening media changer device.
- Improves the performance of various hwmgr commands on large AlphaServer ES80 and GS1280 systems.
- Fixes a "vl\_unwire: page is not wired" panic when the system is configured with gh\_chunks or rad\_gh\_regions enabled and new\_wire\_method = 1.
- Retires the "new\_wire\_method" sysconfig tunable parameter. See "[new\\_wire\\_method Tunable Attribute Retired](#)".
- Corrects the cause of the "pmap\_ssm\_lock\_bit timeout" panic.
- Fixes a problem in which a disk device is inaccessible after physical removal and reinstallation.
- Fixes the cause of a rare hang on an AlphaServer ES47 system that occurs when booting in graphics console mode.
- fixes a problem whereby multiple LSM kernel-initiated plex detaches processed simultaneously by different cluster members across different diskgroups can cause I/O requests to those volumes to hang.

## Patch 27036.00

### *OSFHWBINCOM540*

- Provides environmental sensor support for Superdome power supply sensors.
- Adds support to get live status information for air movers and power supplies on AlphaServer GS1280 systems and to log intrusion packets to the error log.
- Adds support for CPU indictment on AlphaServer GS1280 systems.
- Fixes a situation where, on a panic, the operating system erroneously reboots instead of halting and fails to take a crash dump.
- Distributes the latest .h files for the mcutil program and the event manager.
- Corrects a problem with hwmgr delete while a SCSI scan is in progress.
- Causes the AlphaServer GS1280 platform to accurately report and handle fatal memory errors if the associated CPU is off line.



- Fixes and improves the mcutil program by correcting how bus resets are handled by the program and by enhancing its error reporting capabilities.
- Corrects an issue in which multiple access to changer devices (for example, via Legato/Networker) could lead to a kernel memory fault in the changer driver.
- Fixes problems with DS15 environmental monitoring.

## Patch 27038.00

### *OSFINCLUDE540*

- Corrects a problem introduced in Tru64 UNIX V5.1B-2 that causes a syntax error when compiling programs using the header file /usr/include/netdb.h and defining `_XOPEN_SOURCE_EXTENDED`.
- Corrects the assert macro definition in /usr/include/assert.h.
- Offers the sticky connection feature for a cluster alias.
- Modifies the getaddinfo to compile properly with the POSIX\_SOURCE compiler flag.
- Fixes various problems in the libc functions getdate(), strptime(), callrpc(), strncasecmp() and fork().
- Fixes a problem in the libnuma function ncreate() and the system header <sgtty.h.
- Fixes various problems in the dbx and object file tools dbx, ostrip, strip, mcs, dis, cord, file, and stdump.
- Corrects a problem on systems running Enhanced Security in which the command edauth -R refuses to write user-profile entries to the root partition.
- Changes the use of the configuration file /etc/svc.conf to /etc/nsswitch.conf to allow netgroup data to be provided from LDAP, rather than only from NIS.
- Adds per-binary big page controls to complement the system-wide tunable attributes.
- Fixes an interoperability problem between the curses.h and esnmp.h header files.
- Installs DECthreads V3.20-049.
- Resolves several DECthreads faults and resolves performance issues with certain Java applications.

## Patch 27039.00

### *OSFINET540*

- Provides changes in the bindconfig application (sysman dns) to support BIND 9 version.
- Addresses a formatting error in the output of the tcpdump -r command whereby timestamps may be displayed as negative numbers.
- Fixes invalid permission and ownership of the /usr/bin/deliver IMAP image.
- Fixes a problem with the TruCluster software product running as a NIS server without Enhanced Security installed.
- Upgrades BIND 8 to BIND 9.
- Fixes a cyradm memory fault for an improper directory entry in the /etc/imapd.conf file.
- Corrects the cause of RPC timeout error messages in the ypwhich -m command.
- Adds support to decode TCP SACK packets in tcpdump.
- Corrects the decoding of NFSv3 sattr structure in tcpdump.
- Corrects a problem that occurs during a DNS server configuration using the SysMan DNS server application in which the user is not informed if the named daemon fails to start.
- Changes the use of the configuration file /etc/svc.conf to /etc/nsswitch.conf to allow netgroup data to be provided from LDAP, rather than only from NIS.
- Modifies the prpasswd and rpc.yppasswdd daemons to properly handle /var/tcb/files on a file system from different from /var.
- Corrects a problem that can cause the /usr/sbin/pop3d server to generate a segmentation fault.

- Corrects a problem in which the `/usr/bin/mailauth -ini` command does not create a usable database.
- Corrects a potential security vulnerability in BIND 8 code that could result in a local or remotely exploited Denial of Service (DoS).  
SSRT3653 - BIND v8 — Severity - High
- Adds support for IEEE 802.1Q (VLAN).
- Makes start-up scripts in `/sbin/init.d` world readable.
- 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 potential BIND security vulnerabilities that may result in buffer overflows, unauthorized access, or denial of service (DoS). These may be in the form of local and remote security domain risks.  
SRT2408 BIND — Severity - High  
SSRT2410 BIND — Severity - High  
SSRT2411 BIND — Severity - High
- Fixes the `mkcdsl` command and updates the NIS start-up script to correctly start NIS on the cluster alias.
- Corrects several potential security vulnerabilities where, under certain circumstances, system integrity may be compromised. These may be in the form of improper file or privilege management.
- Fixes a problem in `tcpdump` that causes it to not filter UDP traffic properly.
- Corrects a potential problem in `screend`.
- Corrects a problem in which logins in TruCluster environments using Enhanced Security can hang on any member other than the one serving `/var` to CFS.
- Corrects a problem on systems running Enhanced Security in which the command `edauth -R` refuses to write user-profile entries to the root partition.
- Corrects a problem that occurs when using C1crypt for password encryption on Enhanced Security systems in which users are unable to change their passwords and see the `passwd` command warning "Password not changed: failed to write protected password entry."
- Changes the use of the configuration file `/etc/svc.conf` to `/etc/nsswitch.conf` to allow netgroup data to be provided from LDAP, rather than only from NIS.
- Adds the `retry` option to `snmp_request`.
- Updates Mobile IPv6 code to be compliant with the latest RFC.
- Enables `tcpdump` to display encapsulated IPv6 packets sent over the IPv6-in-IPv6 tunnel.
- Corrects a problem in which a reverse zone lookup entry is not added to the `/etc/namedb/named.conf` file when configuring a cluster node as a DNS server using the `bindconfig` application.
- Resolves a problem in which `tcpdump` can fail with the message "tcpdump: no VLAN support for data link type 0xa."
- Corrects several potential security vulnerabilities where, under certain circumstances, system integrity may be compromised. These may be in the form of improper file or privilege management or remote code execution.

## Patch 27045.00

### OSFKTOOLS540

- Enhances the `inpcb kdbx` extension to display additional PID information.
- Updates the `kdbx audit` extension to accommodate a new device driver state and optionally fetch data appropriately.
- Enhances the `kdbx netstat` extension.

- Fixes a premature termination of the ofile kdbx extension and warning messages in various kdbx extensions.
- Fixes problems in the kdbx u and vnode extensions.
- Provides enhanced kdbx debugging features to include a -A flag for route and to keep inpcb from truncating port numbers.
- Updates Mobile IPv6 code to be compliant with the latest RFC.

#### Patch 27046.00

##### *OSFLAT540*

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

#### Patch 27047.00

##### *OSFLDBBASE540*

- Provides Version V69 of the ladebug debugger to cause it to properly return from a fork call when CATCHFORKS is set.

#### Patch 27048.00

##### *OSFLDBDOC540*

- Provides Version V69 of the ladebug debugger to cause it to properly return from a fork call when CATCHFORKS is set.

#### Patch 27049.00

##### *OSFLEARN540*

- 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.

#### Patch 27050.00

##### *OSFLIBA540*

- Updates the audit system to display additional information for the numa\_syscalls and msfs\_syscall system calls.
- Fixes problems in which acl\_set\_file() fails and returns errno = 22 and does not fail if a file does not exist.
- Corrects a problem in the arena memory allocator that stems from a clash between libc and libnuma on handling errors from the numa\_syscalls system call.
- Provides type checking in EvmVarGet wrapper functions.
- Enhances the AdvFS rmvol utility to allow multiple volumes to be removed with one command.
- Prevents addvol from adding invalid disks into a domain.
- Fixes a problem in XTI, caused by a blocked mutex lock, in which any thread attempting to send an abortive disconnect would hang.
- Installs DECthreads V3.20-029c to fix a problem with floating point data inconsistencies in threaded applications.
- Provides the correct labels for mach events to the audit subsystem.
- Fixes various problems in the libc functions getdate(), strptime(), callrpc(), strncasecmp() and fork().
- Fixes a problem in the libnuma function ncreate() and the system header <sgtty.h>.
- 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.

- Installs version V2.1-120 of /usr/lib/libots3.a and /usr/shlib/libots3.so to fix a problem where long running OpenMP applications might overflow an internal libots3 counter, resulting in a breakdown of thread synchronization.
- Installs DECthreads V3.20-033, which addresses the possibility of floating point errors in threaded programs.
- Corrects the behavior of munlockall in the realtime library (librt).
- Improves null partition checking code.
- Fixes a librt memory leak that can occur when multiple message queue files are opened and then closed. (The memory would be recovered when the process terminates.)
- Expands libpset APIs to enable the caller to get processor set information.
- Installs DECthreads V3.20-029, which fixes problems that may affect threaded programs. DECthreads V3.20-029 is the initial support version of the HP POSIX Threads Library.
- Installs DECthreads V3.20-049.
- Adds support for NEW\_OPEN\_MAX\_SYSTEM (64K) file descriptors to libaio.
- Installs DECthreads V3.20-049a, which fixes a problem that could cause some threaded applications to hang.
- Resolves several DECthreads faults and resolves performance issues with certain Java applications.

## Patch 27051.00

### *OSFLSMBASE540*

- Fixes a problem with vold logging (vold -k -x log) in which the default log file cannot be created.
- Improves LSM's volunmigrate disk processing to produce more informative error messages on command line errors rather than defaulting to a general usage statement.
- Corrects volunmigrate to allow for the full pathname to the device as input and to properly reject partitions with an invalid partition length.
- Fixes an issue whereby the volevac command can appear hung when the number of volumes exceeds 512.
- Provides enhancements to LSM hot-sparing.
- Fixes a problem of adding LUNs off a SWXCR controller to LSM.
- Fixes an issue seen in CLSM that causes a diskgroup import to fail when clsm collision proposals are seen in messages file.
- Fixes a LSM diskgroup deport problem.
- Prevents a "Too Many Volumes" error when attempting to add a volume and the maximum volumes allowed has not been reached.
- Fixes an issue with the initialization of large disks.
- Fixes an issue seen when using KZPAC/SWXCR disks in LSM.
- Corrects several label strings in the SysMan LSM application.
- Fixes problems with SysMan LSM administrator that deal with disk names in "Add disk to LSM" window and the display of disk names in Migrate.
- Allows the LSM voldisk online and voldisk -a online commands to work cluster-wide for devices that have not been replaced. This is useful for rescanning devices that had a temporary disk failure.
- Corrects several potential security vulnerabilities where, under certain circumstances, system integrity may be compromised. These may be in the form of improper file access.

- Corrects a 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.
- Prevents inconsistent LSM volumes when the name of a partition that is being encapsulated matches the name of a current LSM volume.
- Corrects a problem in which smsd triggers LSM configuration errors when querying LSM in a cluster. This fix 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.
- Corrects a problem during the creation of a new plex in which the volmake command reports that associating a subdisk with the plex would cause an overlap with another subdisk when they should not be overlapping.
- Prevents vold from core dumping when attempting to delete a disk that was not initialized properly.
- 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.
- Corrects an I/O performance issue that occurs when CLSM is configured and one member of a cluster goes down unexpectedly.
- Fixes a slow boot when booting several cluster nodes at the same time and CLSM is configured.
- Enables volassist to create a mirror of a striped volume with the mirror having a layout of concat.
- 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.
- Corrects an issue of a cluster node hanging during a boot while another member recovers the cluster root file systems.
- Allow volsave and volrestore to save nconfig/nlog policies for disk groups and restore them appropriately.
- Corrects awk errors for invalid quit statements.
- Corrects a problem in which lsmbstartup does not ignore comment lines in /etc/fstab, thereby making it possible for it to attempt to process that line if it appears to be an entry for the root file system.
- Corrects an internationalization problem in which an instructional message that is output by the LSM voldiskadm command's list option is incorrect because it had not been entered properly in the message catalog.
- Improves null partition checking code.
- Modifies the volassist command to correctly output a warning message for its shrink operations.
- Enables a full recovery to occur in the event of an invalid magic number found in the recovery map.
- Corrects a problem in which the volplex command does not send a disk group option to the usage type utility when both the usage type and disk group options are specified on the command line.
- Allows a mirrored swap plex to be detached when it contains a phantom subdisk.
- Corrects a problem in which vold may dump core when connectivity to a disk group is lost, for example its underlying storage is local to one cluster member and that cluster member leaves the cluster.
- Modifies volsave and volrestore to let them handle public and private offsets of an LSM disk.

- Causes LSM to issue an error if a nonroot user tries to run the volencap command.
- Decreases CLSM boot times in large cluster configurations.
- Fixes problems in vol\_new\_disk ioctl when errors are encountered when loading a disk in the kernel.
- Fixes a problem with the disk group loglen getting set to 0 when adding a disk to a non-rootdg disk group.
- Fixes an LSM problem in which device failures may not be handled properly in a cluster. This is with respect to devices discovered from a previous configuration. The result could be a diskgroup deport.
- Fixes a vold threads problem in which LSM devices may be incorrectly recorded as disk clones where LSM starts up,
- Provides a new LSM EVM event, which is posted when LSM is processing clusterwide plex detaches.
- Modifies LSM to ensure that AdvFS volume links are not required to be named the same as the device they point to.
- Corrects a problem in which the dsfmgr command does not notify LSM of certain disk changes. With this fix, the procedure for replacing a failed boot disk under LSM control should work as described in Section 6.4.6 of the *Logical Storage Manager* manual without generating the error message "lsm:voldisk: ERROR: Device dskxf: define failed."
- Corrects a condition that causes the incorrect use of the volmigrate and volunmigrate commands to return success instead of failure.
- Fix a problem in voldisk modddb that can cause a disk group to become unusable.
- Enhances CLSM performance during boot and disk scanning.
- Fixes a problem in voldisk modddb that can cause a disk group to become unusable.
- Fixes the voldiskadm utility to correctly remove a disk for replacement.
- Fixes a problem whereby LSM mirrored volumes are erroneously left in the SYNC state after a cluster member failure and simultaneous underlying disk storage failure.
- Fixes an error in dissociating swapvol/rootvol plexes.
- Corrects incorrect error messages generated by the voldctl command. With this fix, if voldctl disable fails, the error message returned will be "disable failed" and if voldctl license fails, the error message will be "license check failed."
- Corrects a problem to ensure the correct calculation of subdisk offsets in a plex when associating subdisks to a plex using the volsd command.
- Corrects a problem with non-rootdg disk groups disappearing.

## Patch 27052.00

### OSFLSMBIN540

- Prevents the following panics on systems using LSM or CLSM with low free memory:  
mcs\_lock: lock already owned by cpu  
mcs\_unlock: current lock not found
- Suppresses the following erroneous console warning message that may be provided when cluster root is under LSM control.  
WARNING: cluster root devices are on private buses!  
LSM does not support such configurations.
- Corrects a highly contentious spin lock in the LSM kernel that occurs when running a high number of I/O operations to an LSM volume.
- Corrects a problem to allow CLSM ioctls to return EINVAL when not in a cluster.
- Prevents the dereferencing of a null pointer in volkiostart() when a DRL object has failed and no mirrors belong to the object.

- Fixes a klog inconsistency during node boot in which it is possible that no nodes in the cluster respond to a klog sync request and the booting node simply clears the disk group klog.
- Decreases CLSM boot times in large cluster configurations.
- Fixes problems in the vol\_new\_disk ioctl when errors are encountered when loading a disk in the kernel.
- Fixes the cause of a panic in LSM voldr1\_commit\_delete that occurs during a transaction abort.
- Corrects the cause of a panic in LSM voldr1\_commit\_delete seen during a transaction abort.
- Adds performance enhancements to help reduce LSM cluster-wide plex-detach processing time.
- Fixes a problem whereby LSM does not properly removing plex detach entries from its kernel change log in a TruCluster environment.
- Fixes an issue with kernel plex detach by attempting to further reduce the number of redundant detach events sent from the kernel to vold.
- Fixes a boot problem with the LSM I/O performance change when the root file system is an LSM volume.
- Adds a configurable fastfail option to the LSM driver for mirrored volume I/O.
- Modifies conditions to allow multiple retries to mount the root file system and provide the ability to adjust the retry period.
- Corrects a locking issue in KCL local lock.
- Fixes an "assert wait" panic caused by automatic LSM error tracing when DRD uses LWC (light weight context) switching and many I/O and plex detaches occur.
- Fixes a problem in which I/O requests hang when multiple LSM kernel-initiated plex detaches are processed simultaneously by different cluster members across different diskgroups.

#### Patch 27052.00

##### *OSFLSMX11540*

- Prevents a process hang occurring when large I/O's are performed on multiple drl-enabled LSM volumes simultaneously on Multi-CPU systems.
- Causes zero length I/O requests to be flagged with the proper error message to the user rather than causing a panic.
- Fixes a "clsm\_config" dlm\_lock deadlock issue.

#### Patch 27053.00

##### *OSFLSMX11540*

- Fixes a problem in the LSMSA GUI where the application throws a Null Pointer Exception while trying to display volume and fileset information on systems having ASU file-on-file mounts.

#### Patch 27054.00

##### *OSFMANOP540*

- Updates the uucp(1), gethostbyaddr(3), getnameinfo(3), sys\_attrs\_io(5), kdbx(8), named(8), and volwatch(1) reference pages.
- Updates the dd(1), ksh(1), vi(1), aio\_return(3), sshd2\_config(4), emx(7), dxshutdown(8), newfs(8), prpasswd(8), sshd2(8), and ypset(8) reference pages.
- Revises the EvmEventPost(3) reference page to document the return value EvmERROR\_CONNECTION\_LOST.
- Revises the nsdispatch(3) reference page.
- Revises the getaddrinfo(3) reference page to document the POSIX\_SOURCE compiler flag.

## Patch 27055.00

### OSFMANOS540

- Updates the following reference pages: uucp(1), gethostbyaddr(3), getnameinfo(3), sys\_attrs\_io(5), kdbx(8), named(8), and volwatch(1).
- Updates the following reference pages: awk(1), chmod(1), cp(1), ex(1), find(1), rm(1), uuencode(1), vi(1), and which(1).
- Provides the dig(1) and host(1) reference pages to support BIND 9.
- Updates the following reference pages: poll(2), wait(2), EvmVarSet(3), sys\_attrs\_cam(5), sys\_attrs\_cfs(5), sys\_attrs\_generic(5), sys\_attrs\_inet(5), sys\_attrs\_vfs(5), and sys\_attrs\_vm(5).
- Adds a new reference page, sys\_attrs\_nfs(5).
- Updates the sys\_attrs\_rdg(5) reference page to document an increase in the maximum number of objects in the RDG endpoint and buffer list.
- Adds the named.conf(5) and rndc.conf(5) reference pages to support the upgrade to BIND Version 9.
- Updates the following reference pages: btcreate(8), btextract(8) collect(8), disklabel(8), envconfig(8), fsdb(8), ftpd(8), hwmgr\_show(8), mountd(8), rmvol(8), and ddr.dbase(4). Also adds the following new reference pages to support BIND 9: dnssec-keygen(8), dnssec-makekeyset(8), dnssec-signkey(8), dnssec-signzone(8), named-checkconf(8), named-checkzone(8), nsupdate(8), rndc(8), and rndc-confgen(8).
- Updates the dd(1), ksh(1), vi(1), aio\_return(3), sshd2\_config(4), emx(7), dxshutdown(8), newfs(8), prpasswd(8), sshd2(8), and ypset(8) reference pages.
- Revises the sysconfig(8) reference page to document the cluster interconnect.
- Revises the tcpdump(8) reference page for the VLAN functionality.
- Revises the ifconfig(8), lan\_config(8), niffconfig(8), ping(8) vlanconfig(8), and vlan(7) reference pages for the VLAN functionality.
- Revises the mt(1) reference page to document three new commands, mt reserve, mt release and mt tur, that allow for tape devices to be reserved (via the scsi-2 reserve command) and released (via the scsi-2 release command).
- Revises the envconfig(8) and envmond(8) reference pages for the environmental monitoring facilities (/usr/sbin/envmond, /usr/sbin/envconfig) to support the AlphaServer GS1280 hardware platform.
- Revises several of the SSH reference pages to address issues and problems with SSH, such as interoperability with other SSH implementations, SSH client/server configuration files compatibility issues, and the lack of IPV6 support.
- Revises the sys\_attrs\_ee(5) reference page to document the new ee subsystem attribute link\_check\_interval.
- Revises the newfs(8) reference page to document the -M command option to newfs, which lets you specify permissions of an MFS root directory when it is first created.
- Revises the vdump(8) reference page to change the statement for the -b option to be a maximum of 2048.
- Adds the new reference page chatr(1).
- Revises the sys\_attrs\_proc(5) reference page to add a new tunable attribute, executable\_data.
- Adds the new reference page javaexecutedata(8).
- Revises the fwupgrade(8) reference page to document a new feature that allows the specified firmware update image to be located on a BOOTP server in a connected network.
- Revises the collect(8) and pmgrd(8) reference pages to document new AdvFS features.
- Revises the codconfig(8) reference page to include new supported platforms.
- Revises the xmesh(1) reference page to include new supported platforms.
- Revises the sys\_attrs\_clubase(5) reference page to document the change to the cluster\_rebuild\_delay.



- Revises the `sys_attrs_lsm(5)` reference page to document scalability and performance with LSM spin locks.
- Revises the `fuser(8)` reference page
- Revises the `binlogd(8)` reference page to document information about adding a CDSL (Context Dependent Symbolic Link) to the binlog archive directory.
- Revises the `voliod(8)` reference page to document a change to the `voliod -f set 0` command.
- Revises the `evminfo(1)` reference page to document the addition of the `auth` option to the `-verify` option, which allows users to check for syntax errors in the EVM authorization file.
- Revises the `kdbx(8)` reference page to document the addition of the `-A` flag to `route`.
- Revises the `netstat(1)`, `sys_attrs_inet(5)`, `sys_attrs_net(5)` and `route(8)` reference pages.
- Revises the `chatr(1)` and `sys_attrs_vm(5)` reference pages.
- Revises the `volrestore(8)` and `volsave(8)` reference pages to document the change which will allow `volsave` and `volrestore` to save `nconfig/nlog` policies for disk groups and restore them appropriately.
- Revises the `volassist(8)` reference page.
- Revises the `ldapcd.conf(4)`, `ldapusers.deny(4)`, `netgroup(4)`, `svc.conf(4)`, `nsswitch.conf(5)`, `edauth(8)`, `nss2svc(8)`, `nssetup(8)`, and `svcsetup(8)` reference pages.
- Revises the `nslookup(8)` reference page.
- Revises the `rcp(1)` reference page.
- Revises the `pmgrd(8)` reference page to correct the path name and the file name for the `pmAdvfs.mib` file.
- Revises the `migrate(8)` reference page to add information about interrupts that can occur when direct I/O is enabled.
- Revises the `xmesh(1)` reference page to add the Xmesh resource file information in the FILES section.
- Revises the `snmp_request(8)` reference page to document the new `-r` option.
- Revises the `sys_attrs_net(5)` reference page to add `pfilt_loopback` and `pfilt_physaddr`.
- Revises the `sys_attrs_ufs(5)` reference page to document the tuning of the UFS attribute `delay_wbuffers` using `sysconfig`.
- Revises the `sys_attrs_alt(5)` reference page.
- Revises the `ip(7)` reference page.
- Revises the `sys_attrs_inet(5)` reference page to update the default value for `ipqmaxlen` to 2048 (instead of 1024).
- Revises the `sys_attrs_bcm(5)` reference page.
- Revises the `chatr(1)` reference page.
- Revises the `audit_tool(8)` reference page to document a change that allows a user to specify a path to archived audit logs that the `audit_tool` will follow for all logs instead of the path recorded in the `audit_log_change` events (which by default is `/var/audit`).
- Revises the `cs(1)`, `sh(1b)`, `ip6rtrd.conf(4)`, `ntp.conf(4)`, `sys_attrs_ipv6(5)`, `sys_attrs_vfs(5)`, `niffmt(7)`, `dump(8)`, `freezefs(8)`, `fwtmp(8)`, `ip6rtrd(8)`, `restore(8)`, and `xntpd(8)` reference pages.
- Revises the `sys_attrs_lsm(5)` and `disklabel(8)` reference pages.
- Revises the `mt(1)`, `netstat(1)`, `sys_attrs_ee(5)`, `sys_attrs_inet(5)`, `sys_attrs_vm(5)`, `collect(8)`, `envconfig(8)`, `fuser(8)`, `ifconfig(8)`, and `kdbx(8)` reference pages.
- Revises the `btcreate(8)` and `psradm(8)` reference pages.

## Patch 27060.00

### *OSFNETCONF540*

- Corrects several potential security vulnerabilities that have been reported on systems using an IPsec configuration of tunnel mode ESP without authentication. Under certain

circumstances, a remote attacker could force an error so that a portion of a plain-text message can be intercepted by the attacker. Corrects the following vulnerability:

SSRT5957 - IPsec (Severity - High)

- Fixes a problem in which a stack trace occurs during quicksetup if the host name and any NFS parameters are changed.
- Modifies the netconfig application to prevent breaking automation while using the SysMan command-line interface.
- Corrects the netconfig application to avoid a stack trace that occurs while configuring a token ring adaptor on a different node.
- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file access.
- Corrects a problem in which when DHCP is selected for a network interface card, netconfig places invalid data in the /etc/hosts file.
- Changes the use of the configuration file /etc/svc.conf to /etc/nsswitch.conf to allow netgroup data to be provided from LDAP, rather than only from NIS.
- Fixes a problem in the SysMan nfs\_export application in which adding a host to the rw-access list does not take effect.
- Fixes a problem in the SysMan nfs\_export application in which an inappropriate message is displayed when a nonroot user runs it.
- Corrects a problem in which the network wizard exits when running as nonroot.
- Corrects a potential security vulnerability in IPsec/IKE (Internet Key Exchange) with Certificates. This potential vulnerability is remotely exploitable, resulting in unauthorized privileged access.

SSRT3674 - IPsec/IKE (Severity - High)

- Fixes a problem in which SysMan route does not handle the destination name input correctly

## Patch 27062.00

### OSFNFS540

- Adds a name and IP address cache to NFS mountd to limit problems seen with DNS timeouts.
- Corrects a problem with pcnfsd that occurs when it is running in a cluster and sends the member's IP address instead of the cluster alias.
- Fixes the warning message in the daemon.log file when automount starts.
- Provides an option that allows users to specify port number for mountd.
- Fixes rpc.lockd to send UNIX authentication.
- Corrects a problem in which autofs skips exported file systems that belong to netgroups in the remote NIS domain, while a mount command for the same exported file system works.
- Fixes problems due to mountd not blocking SIGHUP when processing NFS exported file systems.
- Fixes multiple defects in AutoFS user space and kernel code.
- Enables mountd to correctly handle entries with multiple lines of input in an exports file.
- Makes start-up scripts in /sbin/init.d world readable.
- Fixes an issue with the rpc.lockd daemon's message passing-style RPCs, where replies are sent to the IP address of the lock's caller\_name field instead of to the call message's source.
- Modifies the mountd daemon to correct a core dump problem.
- Restore exports file (-root=hostlist) behavior.
- Changes the use of the configuration file /etc/svc.conf to /etc/nsswitch.conf to allow netgroup data to be provided from LDAP, rather than only from NIS.
- Enables AutoFS to handle loopback mounts correctly, specifically regarding failed attempts to use AutoFS to access a loopback mount via a cluster alias.

- Fixes the behavior of the rpc.rquotad daemon in a cluster.
- Addresses a mountd problem in which spurious signals can cause mountd to dump core and a second problem that causes .INCLUDE parsing to function incorrectly.
- Enables AutoFS mount-on paths to have the correct maximum length when AutoFs files are mapped directly.
- Fixes an rpc.lockd problem involving lock contention, in which lock requests accumulate on a linked list that is searched with each kernel lock manager poll and result in excessive CPU consumption
- Fixes an issue in which an error was being reported with an empty direct map.
- Fixes a problem where rpc.statd does not correctly create the sm and sm.bak directories on startup, causing rpc.statd to exit.

#### Patch 27064.00

##### *OSFOBSOLETE540*

- Updates Tru64 UNIX version information in accordance with HP rebranding efforts.

#### Patch 27065.00

- Updates Tru64 UNIX version information in accordance with HP rebranding efforts.
- 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.

##### *OSFOEMBASE540*

#### Patch 27068.00

##### *OSFPGMR540*

- Fixes a sendmail registration issue with PSM.
- Updates sysconfig to use the cluster interconnect, thereby allowing for a greater SSI collaboration. This will help with changing variables on hung systems, single user systems, and normal running systems.
- Corrects several potential security vulnerability where, under certain circumstances, system integrity may be compromised. These may be in the form of improper file or privilege management.
- Corrects a core dump problem that occurs when output of the lint command for nonexistent file is supplied to error.
- Fixes many small problems in dsfmgr.
- Fixes various problems in the dbx and object file tools dbx, ostrip, strip, mcs, dis, cord, file, and stdump.
- 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.
- Corrects a problem that occurs when stdin, stdout and stderr are closed and cfg\_connect fails to connect to cfgmgr.

#### Patch 27069.00

##### *OSFPRINT540*

- Provides LPD support for the following HP Printers:
  - Color LaserJet 4650
  - Color LaserJet 9500

- LaserJet 4200
- LaserJet 4250
- LaserJet 4300
- LaserJet 4350
- LaserJet 9040
- LaserJet 9050
- Fixes a problem with `lpstat -o` in handling hyphens in printer names. If there is a hyphen in printer name, `lpstat -o` was interpreting it as a job-id instead of printer name.
- Corrects a problem that exists with print jobs submitted on clusters that occurs when an "on" attribute in `/etc/printcap` is specified for one node and a job is submitted on another node.
- Modifies `lpd` to fix `/etc/hosts.lpd` case sensitivity; for example, "node.domain" treated the same as "Node.Domain."
- Makes start-up scripts in `/sbin/init.d` world readable.
- Improves printing (`lpd`) maintainability, including the addition of new `DEBUG lpr.log` messages, a `DEBUG` message screening feature, and an improved remote printing connection-retry scheme.
- Fixes a problem in which job pages are missing when the `jj` printcap flag is set to 1.
- Fixes a problem in which temporary files are left in root (`/`) when `lpd` gets bad file names from `/etc/printcap`.
- Provides new delay timers to address timing problems.
- Provides a new `syslog DEBUG` message for `lpr.log` to debug bad remote job file names.
- Fixes a problem that causes queues that use the `lpf` output filter (and potentially other output filters) to hang when the print subsystem is used heavily; for example, 300 queues used continuously.
- Fixes an error in which remote job reprinting does not occur when needed.
- Adds a new `/etc/printcap` option, `sr`, to suppress job reprinting.

#### Patch 27070.00

##### *OSFRCS540*

- 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.

#### Patch 27071.00

##### *OSFRIS540*

- Corrects a joint failover problem and an extracted RIS area issue.

#### Patch 27072.00

##### *OSFSCCS540*

- 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.

## Patch 27073.00

### *OSFSDE540*

- Fixes various problems in the dbx and object file tools: dbx, ostrip, strip, mcs, dis, cord, file, and stdump.
- Fixes the way the spike code optimization tool handles -arch and -tune options. Using these options with previous versions of spike resulted in unproductive error messages and/or system crashes.
- 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.
- Fixes an internal error with the /usr/bin/spike optimization command.
- Modifies the cflow utility to store its temporary files in a temporary directory created using the /usr/bin/mktemp utility. If cflow cannot create a temporary directory it terminates with an exit value of 1.
- Fixes a fatal error in /usr/bin/spike.
- Fixes a problem in which the prof -pixie -testcoverage <exe <exe.Counts sometimes reports invalid source line number ranges.
- Changes the use of the configuration file /etc/svc.conf to /etc/nsswitch.conf to allow netgroup data to be provided from LDAP, rather than only from NIS.
- Fixes performance tool failures on Sierra Clusters Parallel File Systems.
- Corrects an inaccurate error message in spike optimizer.

## Patch 27075.00

### *OSFSER540*

- Fixes a panic on a system with an Oxygen VX1 graphics card when the X server is killed while it is starting up.
- Fixes a problem in which the X server shared memory extension sometimes does not display images properly in windows of depth 16 or 24.
- Corrects a problem that occurs with the Oxygen VX1 graphics card in which XCopyPlane copies all bitplanes rather than only the requested bitplane.
- Fixes a memory leak in the X server Panoramix/Xinerama Extension that can cause a process to core dump.
- Fixes a problem in which the X server's command line option to turn off VESA Display Power Management Signalling (-dpms) does not work.
- Corrects a problem that can cause the X server to hang every 49 days on systems with PowerStorm 4D40T, 4D50T, 4D51T, or 4D60T graphics options.
- Fixes a problem in which the X server may crash if the RADEON layered product kit is installed, 3D rendering is disabled, and a client queries the GLX extension.
- Corrects a potential file permissions vulnerability and a potential buffer overflow in the X Window System. The potential vulnerabilities are locally exploitable, resulting in unauthorized privileged access.
- Fixes a problem with the X server when the Panoramix extension is enabled in which clients running on Linux® systems with their displays directed to the Tru64 UNIX X server will fail to start.
- Fixes a problem in the X server that causes mouse buttons to stop working as expected.

## Patch 27077.00

### *OSFSERVICETOOLS540*

- Corrects a problem that can cause the collect utility to die from a memory access fault on systems with a large number of AdvFS domains and filesets.
- Provides the collect utility with support for long disk names.

- Corrects a potential security vulnerability that could result in non-privileged users gaining unauthorized privileged access on the system.
- Provides the ability for the collect utility to salvage data files with missing termination records.
- Provides the collect utility with the ability to report local and remote I/O access statistics for cluster storage devices in a TruCluster Server environment.
- Corrects a problem of compromised AdvFS data integrity which occurs when a new AdvFS domain or fileset is added while collect is running.
- Updates the collect utility to Version 2.0.5.
- Provides new collect features, including AdvFS monitoring and CPU and memory metrics on a per RAD basis.
- Fixes floating-point exception in collect.
- Corrects a condition in which collect cannot create a new data file from a data file that does not include a termination record.
- Corrects a condition in which the concatenation of two or more of collect's data files prevents access to last record and collect generates an "out-of-sync" error. With this patch, collect will print out the last record when the -l option is used on playback.
- Allows collect to gather AdvFS statistics when the number of active AdvFS filesets exceeds 128.
- Causes collect to exit with the following message if it is run with the -H option and the directory linking to /var/adm/collocated does not exist:  
chmod or chown failed.: No such file or directory
- Corrects a problem in which collect hangs on a AS8200 server and CPU resources would rise to 99.9% until a Ctrl/c is issued.
- Fixes the problem where collect displays wrong values of usrtim and systim.
- Cleans up resource usage caused by collect.
- Fixes an issue with the collect on NUMA platforms in which collect RAD CPU Summary statistics are inconsistent with the single statistics.
- Corrects a problem in which collect stops working intermittently when data collection is done with compression.
- Fixes a memory leak problem with collect on NUMA systems.
- Fixes a problem with the collect that cause it to report incorrect packet count information when a large number of incoming/outgoing packets occur per second.
- Improves collect so it can properly handle dynamic changes to resources such as AdvFS file systems and disks. The collect utility now also ensures that errors are logged through syslog so they are not lost when it is run in historic mode.
- Corrects a condition that causes collect to display the error "Mini-Header out of sync" when playing back data files.
- Improves input validation and the handling of process lists supplied via the collect -P command.
- Corrects a problem in which memory reports generated by collect for a system on NUMA platform are different from those generated on a non-NUMA system.
- Corrects a problem with collect when it is used with the -l option to seek to the last record of a data file whose size is less than 312 bytes.
- Corrects a problem with collect incorrectly reporting RAD and memory statistics on systems where the RADs are not serially numbered.
- Corrects a problem of collect stopping after 48 hours, when running in historical mode.

## Patch 27079.00

### OSFSYSMAN540

- Fixes a smsd core dump problem in a cluster.
- Fixes repeated AdvFS domain panics triggered by smsd.
- Fixes a problem in dxaccounts, where the Modify option dumps core when more than one NIS user is selected on the NIS client.
- Corrects a problem in the SysMan advfsmgr utility, where the size of a very large fileset is shown as zero.
- Fixes a problem where sysman -station does not work in the French locale.
- Fixes a problem in which a stack trace occurs during quicksetup if the host name and any NFS parameters are changed.
- Fixes a problem in which SysMan ppp\_options does not show the ttyname and speed values correctly.
- Fixes a problem where "sysman users" does not allow root to change the password.
- Corrects a problem in the function of SysMan ufsmgr "expand."
- Fixes a condition that can cause a copy-and-paste of a NIS-group to hang.
- Modifies the netconfig application to prevent breaking automation while using the SysMan command-line interface.
- Fixes problems with the dxaccounts application on systems with ASU installed.
- Corrects a problem that occurs with ntpconfig that causes the error message "keyword authdelay unknown" to be written into the daemon.log file whenever XNTPD V4 starts.
- Fixes a problem in which a stack trace occurs in the ntpconfig application during attempts to delete duplicate entries.
- Fixes a problem in which dxaccounts deletes the home directory upon modification
- Changes the incorrect choice LSMnopriv to LSMnoprv in the Usage: button option of the diskconfig application.
- Corrects a problem in the SysMan account management application in which an invalid error message is displayed while deleting a user who does not have a home directory.
- Enables the multi-column List widget implemented in SUIT to use any value specified for the -height option in the curses domain.
- Corrects several potential security vulnerabilities where, under certain circumstances, system integrity may be compromised. These may be in the form of improper file access.
- Corrects the cause of a core dump that occurs when the usermod command is used with the -x pc\_synchronize option.
- Corrects a problem that causes the userdel command to core dump when the shell field is empty in the passwd file.
- Corrects a problem that causes the usermod command to not work as expected with NIS +/- users.
- Fixes minor problems with the account management tools.
- 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.
- Updates the account management tools to use the latest versions of the ASU (Advanced Server for UNIX) API calls when ASU is in use on the server.
- Provides support for SmartArray disk controllers. Without this fix, if the SmartArray product is installed on the system the SysMan Station hardware view will fail to operate.
- Modifies the useradd command to correctly manage default and template data properly. The problem showed up most notably with the useradd -p command would producing the message "Password must be between 32 and 80 characters."

- Corrects a problem in some cluster configurations that causes the SysMan Station to fail to generate the hardware view and a Java stack trace is generated indicating that the routine "HardwareLayout.fancyPlace" was being executed at the time of the trace.
- Provides enhancements for file system suitlets.
- Corrects a problem in which the SysMan Station display is not updated when a component (for example CPU) registration and deregistration events occurs.
- Corrects a problem in which if the Reload/Refresh button is pressed during the running of any of the SysMan tools in a browser, the tool hangs, thereby requiring the user to restart the browser to continue with the tool.
- Corrects a condition in which the daemons used by Insight Manager fail and do not restart when restarting the network using the SysMan Menu.
- Corrects a problem in which when SysMan is used to change the configuration of a network adapter card, it would also reconfigure other cards, thereby causing a cluster connectivity interruption
- Prevents users from attempting to configure a network adapter card that is part of a NetRAIN set.
- Updates the SysMan Menu layout in online help.
- Fixes a problem that occurs when dealing with white spaces in a "Filter by Comment" search of the SysMan Account Management application.
- Corrects a problem that occurs when DHCP is selected for a network interface card, the netconfig command places invalid data in the /etc/hosts file.
- Fixes a problem encountered when btape TCL scripts are executed by nonprivileged users.
- Adds \$quote directive to the message catalog.
- Corrects a confusing error message seen when ASU is installed and a user runs useradd on a non-PDC server.
- Fixes potential issues with system security during the creation of temporary files.
- Corrects a problem in which the SUIT multicolumn list incorrectly handles the programatic sizing of the columns, thereby causing data to be truncated.
- Corrects a problem in which SysMan Station fails when the HOME environment variable is not set or points to a nonexistent directory.
- Corrects a problem in which the SysMan Station does not allow the relocation of CAA resources in the CAA view by dragging and dropping resource objects.
- Corrects a condition that causes the useradd command to create the same UIDvalue for a new user when the /etc/passwd file contains the expression #?.
- Corrects a problem on systems with multipath SCSI devices in which the SysMan Station Hardware View incorrectly displays the hardware components and their relationships.
- Fixes a problem with the display of the appropriate title label in the Manage Local Users and Manage NIS Users SysMan applications in a Java domain.
- Fixes a problem with the dialogue box that appears when remote printing and Advanced Server printer types are selected during the Add Printer configuration.
- Improves the input validation for Client, Server and IP address fields for the pap-secrets and chap-secrets file.
- Corrects a problem in curses mode in which AlphaServer Management Server tool incorrectly displays selected nodes and loses column headers after viewing help.
- Corrects a condition that causes double clicking table headers in Manage CPUs and AMS causes traceback.
- Corrects a class loader hang that occurs when run by non-root users.
- Corrects a condition that causes the setld SUITlet to fail when it is run from a browser or SysMan Station.
- Corrects a userdel -D failure to a remove user from the /etc/passwd file on Enhanced Security systems.



- Eliminates two warning messages from the SysMan Security Configuration tool. One appears when changing the root password on systems with ASU installed and the other appears when a user tries to trim auth logs before they are created.
- Corrects a problem in which running the usermod command with certain options, such as -x distributed grace\_limit, results in a “Cross-device link” error when the var and yp directories belong to two different filesets, on a NIS master.
- Corrects a problem in which when a node in a cluster is rebooted, the SysMan Station daemons running on each of the other nodes core dumps and a core file can be seen on the root (/) directory.
- Changes the use of the configuration file /etc/svc.conf to /etc/nsswitch.conf to allow netgroup data to be provided from LDAP rather than only from NIS.
- Corrects a problem that occurs when new hardware objects are inserted on a node in a cluster and smsd on each node displays the newly added hardware object as being added to the local node.
- Fixes a problem in which ntpconfig posts two events for each modification of the ntp.keys file.
- Corrects problems in which SysMan drdmgr dumps tcl stack when a user tries to manage devices or file systems of a cluster node that is down.
- Corrects a condition that causes failures when locking and unlocking multiple selected users.
- Modifies netconfig to avoid tcl stack dumps that occur when a user who has no HOME directory tries to view the network daemon status of a cluster alias or a remote node or tries to set up network interface cards on a remote node.
- Fixes a problem with the validation of the account expiration date in usermod.
- Corrects a problem in which the network wizard exits when running as nonroot.  
Corrects a problem in which the Modify and Delete keys are disabled when using the arrow keys to move through a list.
- Modifies SysMan SUIT so that suitlets such as Manage CPUs and Event Viewer can successfully be launched in JAVA display mode.
- Fixes a java null pointer exception problem when starting a sms client whose initial view window has been deleted.
- Fixes the cause of a core dump that occurs during a copy/paste operation of a Local/NIS template
- Modifies the useradd and usermod commands to validate that the shell argument passed is a valid executable file.
- Improves the input validation in SysMan pppconfig.
- Fixes a problem with the copy/paste operation of dxaccounts.
- Fixes a problem with the creation of new user using useradd.
- Modifies the SysMan Display Mounted File Systems application to correct a problem with file system size.
- Modifies the SysMan Display Currently Mounted File Systems application to correct a problem in which it incorrectly lists file-on-file mounts.
- Modifies the SysMan Mount File Systems application to correct a problem in which it incorrectly mounts entries in /etc/fstab file.
- Fixes a problem where the SysMan Unmount File Systems application does not display file-on-file mounts correctly.
- Addresses a problem with the usermod command when it is used to perform operations on NIS local plus or minus users.
- Modifies the dop command and SysMan dop to correct various problem.

- Addresses a problem with the usermod command when it is used to move a user's home directory to a subdirectory of the old home directory, causing the deletion of the old home directory (and thereby, the newly created home directory).
- Addresses a problem with the userdel command when it is used to delete an NIS local plus or minus user.
- Fixes a problem of not being able to change the user login name in SysMan account management application.
- Fixes a problem with dxaccounts with C2 security enabled, where dxaccounts cannot set the null password for a user even if the user's database entry allows the user to have a null password.
- Corrects a problem in which diskconfig does not delete old alias name when a new alias name is specified.
- Fixes the cause of a core dump that occurs with in dxaccounts when the Delete operation is attempted on a default template in Local Templates view under enhanced security.
- Fixes a problem with SysMan Menu when it is invoked through a browser.
- Corrects a problem where SysMan Station does not show the long name of CAA resources.
- Corrects a problem in which using the usermod -d -m command to specify the creation of a new home directory fails if the new directory contains a trailing slash (/).
- Corrects a problem in which diskconfig does not display UFS file system creation errors.
- Modifies the Quicksetup application for evaluating Printer Name and Printer Server data.

#### Patch 27080.00

##### *OSFTCLBASE540*

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

#### Patch 27083.00

##### *OSFTRUETYPE540*

- Corrects a potential file permissions vulnerability and a potential buffer overflow in the X Window System. The potential vulnerabilities are locally exploitable, resulting in unauthorized privileged access.
- Corrects a potential file permissions vulnerability and a potential buffer overflow in the X Window System. The potential vulnerabilities are locally exploitable, resulting in unauthorized privileged access.

#### Patch 27084.00

##### *OSFUUCP540*

- Corrects a problem with the uucp command in which it changes the time stamps of the destination file when the destination file already exists as a directory.
- Enables the tip command to log into the member-specific log file.
- Corrects the path for the aculog file and gives the file appropriated permissions.
- Makes start up scripts in /sbin/init.d world readable.
- Corrects several potential security vulnerabilities where, under certain circumstances, system integrity may be compromised. These may be in the form of improper file or privilege management.
- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised when a buffer overflow occurs in the uucp utility. Buffer overflows are sometimes exploited in an attempt to subvert the function of a privileged

program and possibly execute commands at the elevated privileges if the program file has the `setuid` privilege.

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

## Patch 27085.00

### *OSFX11540*

- Corrects a problem that occurs when attempting to open a Java applet on a V1.7 Mozilla browser displayed from a Redhat Linux RHEL 3 (update 6) to a Tru64 UNIX V5.1B-3 graphics console. The browser closes and `dtwm` spins at 99% CPU usage.
- Resolves security vulnerabilities in X `Pixmap` routines used in the Motif library.
- Modifies the online help description for `Togglekeys` in `accessx`.
- 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.
- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file access.
- Corrects a problem, where, under certain circumstances, the `XmCvtXmStringToCT()` function does not correctly convert a compound string to a string in compound text format.
- Fixes a problem where a Chinese character whose byte sequence contains `0x9b` cannot be entered with `dxhanziim` or cut and pasted.
- Provides an updated keyboard map for the Russian 3R-LKQ48-BT keyboard model.
- Fixes a display width mismatch problem in the `zh_CN.GB18030` locale.
- Fixes a problem with `xterm` while displaying a compound text that is converted by `XmCvtCTToXmString()`.
- Modifies `XmbTextListToTextProperty()` and `XmbTextPropertyToTextList()` to support 4-byte length UTF-8 characters in the Compound Text handling.
- 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 various problems with the X font server and with the X server's interaction with X font servers.
- Prevents application failures when an application specifies very large timeout values to X Toolkit library (`Xt`) routines.
- Resolves a drag and drop problem across the screen in multi-head systems.
- Corrects a potential file permissions vulnerability and a potential buffer overflow in the X Window System. The potential vulnerabilities are locally exploitable, resulting in unauthorized privileged access.
- Resolves a potential buffer overflow within the X `Pixmap` routines.

## Patch 27086.00

### *OSFXADMIN540*

- Corrects a potential security vulnerability in the XDM (X Display Manager) software. This potential vulnerability, which may be locally and remotely exploitable, could result in a denial of service (DOS), unauthorized privileged access, or both.
- Corrects a problem in which host icons overlap in the `dxhosts` application.

- 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.
- Corrects the way the dxkerneltuner application handles memory.
- Fixes several potential security vulnerabilities where, under certain circumstances, system integrity may be compromised. These may be in the form of improper file access.
- Corrects a problem in which the dxkerneltuner application when invoked with the -power option does not remove the temporary files in /tmp when the application is closed.
- Corrects a problem in which dxproctuner does not display data in the fr\_FR local.
- Corrects a problem in which the dxproctuner application, when invoked, either gives error message or dumps core and does not work.
- Corrects a problem in which a temporary file is not deleted when dxarchiver is closed.
- Corrects a core dump problem in the dxfileshare application when the /etc/exports file is not present in the system and a user tries to add a new entry.
- Corrects a problem in which the dxkerneltuner application dumps core when cancel is selected in the mmsess subsystem.
- Corrects a problem in which the dxpower application shows an error when the Spin Down Disks option is selected.
- Fixes problem with dxproctuner in which the toggle menu item View Processes for All Users does not filter the output for non-root users.

#### Patch 27089.00

##### *OSFXDEMOS540*

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

#### Patch 27094.00

##### *OSFXLIBA540*

- Corrects a problem that occurs when attempting to open a Java applet on a V1.7 Mozilla browser displayed from a Redhat Linux RHEL 3 (update 6) to a Tru64 UNIX V5.1B-3 graphics console. The browser closes and dtwm spins at 99% CPU usage.
- Corrects a problem, where, under certain circumstances, the XmCvtXmStringToCT() function does not correctly convert a compound string to a string in compound text format.
- Fixes a problem where a Chinese character whose byte sequence contains 0x9b cannot be entered with dxhanziim or cut and pasted.
- Modifies XmbTextListToTextProperty() and XmbTextPropertyToTextList() to support 4-byte length UTF-8 characters in the Compound Text handling.
- 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.
- Prevents application failures when an application specifies very large timeout values to X Toolkit library (Xt) routines.
- Corrects a potential file permissions vulnerability and a potential buffer overflow in the X Window System. The potential vulnerabilities are locally exploitable, resulting in unauthorized privileged access.

## Patch 27095.00

### *OSFXMIT540*

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file access.
- Fixes various problems with the X font server and with the X server's interaction with X font servers.
- Corrects a potential file permissions vulnerability and a potential buffer overflow in the X Window System. The potential vulnerabilities are locally exploitable, resulting in unauthorized privileged access.

## Patch 27100.00

### *OSFXSYSMAN540*

- Fixes a problem in dxaccounts in which the Modify option dumps core under certain circumstances.
- Fixes several problems with dxaccounts on a system with ASU installed.
- Fixes several potential security vulnerabilities where, under certain circumstances, system integrity may be compromised. These may be in the form of improper file access.
- Corrects a problem with the userdel command that can cause it to dump core when the shell field is empty in the passwd file.
- Corrects a problem with the usermod command not working as expected with NIS +/- users.
- Fixes several minor problems with the account management tools.
- A potential security vulnerability has been discovered, where under certain circumstances, system integrity may be compromised. This may be in the form of passwords that have a length outside of the intended range.
- Updates the account management tools to use the latest versions of the ASU (Advanced Server for UNIX) API calls when ASU is in use on the server.
- Updates the SysMan Menu layout in the online help.
- Corrects a problem in which extra leading and trailing spaces from the FIND dialog box in dxaccounts are not trimmed before being used, which results in a search failure.
- Modifies dxaccounts to display the account expiration date at first view.
- Fixes a condition in which a disabled I/O port incorrectly reports 100% utilization.
- Causes the display of instructions for creating the /dev/marvel\_pfm file when the file is not present on the system.
- Fixes a condition in which the menu does not wrap around when the window width is too small to accommodate all the menu items.
- Adds the xmesh resource file /usr/lib/X11/app-defaults/Xmesh to customize xmesh.
- Allows the adjustment of the color key to change the percentage range for each color.
- Modifies the dxaccounts application to prevent core dumps under the following circumstances:
  - When a UNIX login name is modified and its corresponding PC account is not remapped to the new login name.
  - When a user tries to modify the UNIX and PC account login names and then tries to see information in PC user view.
  - When the root user is selected from the drop-down list for modification.
  - During a copy-and-paste operation of the Local/NIS template
  - In the Retire function.
- Fixes a problem where, under certain circumstances, the Display Window application (dxdw) displays garbled characters in the transcript area of the application.

- Modifies dxaccounts to allow it to change the local user password in Enhanced Security environment.
- Corrects a problem in which the locking and unlocking of multiple selected users did not happen under C2 security.
- Fixes the cause of a core dump that occurs during copy/paste operations of Local/NIS template.
- Corrects a problem in which the dxaccounts application incorrectly creates a default expiration date of Feb 1, 1995 for new user accounts.
- Fixes a problem with copy/paste operations using dxaccounts.
- Fixes a problem that occurs when creating a new user with useradd.

#### Patch 27102.00

##### *OSFCDSABASE540*

- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of compromised private RSA keys.
- Corrects a problem in which during a CDSA configuration the mod\_install program core dumps.

#### Patch 27106.00

##### *OSFIPSECBASE540*

- Corrects multiple potential vulnerabilities identified on HP Tru64 UNIX operating systems running IPsec, which uses the Internet Security Association and Key Management Protocol (ISAKMP). The vulnerabilities could be exploited remotely to cause Denial of Service (DoS).
- Corrects a potential security vulnerability in IPsec.

#### Patch 27107.00

##### *OSFLDPAUTH540*

- Allows the ldapd to exit when there is a permanent error. When ldapd exits, init respawns the daemon allowing ldapd to recover and function normally.
- Modifies ldapd to prevent it from crashing under the following conditions:
  - When resolving group codes with very large GIDs.
  - When the LDAP Directory Server is unavailable.
- Corrects a condition in which the login process may crash when LDAP users, or users belonging to an LDAP group, attempt to log in.
- Fixes a problem with ldapd that prevents LDAP users from being authenticated, even when they are providing the correct password.
- Changes the use of the configuration file /etc/svc.conf to /etc/nsswitch.conf to allow netgroup data to be provided from LDAP, rather than only from NIS.
- Fixes problems caused when the UID/GID of an Active Directory user is zero.
- Corrects a condition in which the su command issues multiple "Sorry" messages when multiple SIA mechanisms are in use; for example, when LDAP is configured for user accounts.
- Ensures that login sessions are properly established for LDAP and BSD accounts when multiple authentication mechanisms are configured.
- Addresses the way that the ldapd handles when the ldap server becomes unavailable.

## Patch 27110.00

### *OSFOPENGL540*

- Fixes a problem in which the OpenGL `glTexCoordPointerEXT()` command can cause a segmentation fault.
- Restores two extensions missing from OpenGL library to the list of GL extensions it supports for indirect rendering.

## Patch 27113.00

### *OSFSSHBASE540*

- Corrects a potential issue with `scp2` and `sftp2`.
- Fixes the cause of memory leaks in `sshd`.
- 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.
- Fixes the error "Xauth data does not match fake data." that can occur when multiple SSH sessions from the same client are open on different cluster member nodes.
- Modifies the `ssh-pubkeymgr` script to change the default keyfile name to `user-host`, and to simplify the procedure for enabling a key for a remote login.
- Fixes a problem in SSH in which when attempting SSH TCP port forwarding the SSH server handling the forwarding would die.
- Corrects a problem that occurs when booting a during a file system full situation in which the `ssh-validate-conf` utility attempts to write to the files `/etc/ssh2/sshd2_config` and `/etc/ssh2/ssh2_config`, thereby causing them to be zeroed out.
- Corrects a potential security vulnerability.
- Corrects a misspelling in the `ssh-hostbased-setup` utility message "is not running a compatible `sshd`, skipping."
- Corrects a problem in which `scp` does not check whether the source and destination were the same file, thereby causing the file to be truncated to zero bytes.
- Corrects a condition in which if a user connects to a cluster, performs two SSH localhost, and then tries to start an X application, an error message of X connection is broken is displayed.
- Corrects a problem in which the `SSH-hostbased-setup` utility does not handle host names containing a hyphen (-).
- Corrects the handling of `chroot` users via `ssh` with Enhanced Security enabled.
- Fixes a problem with `scp` where, in some cases, the source file could be cleared.
- Fixes an issue with SSH V3.2.3 host-based authentication when using the `MapFile` configuration option.

## Patch 27114.00

### *OSFSSOSSL540*

- Corrects a potential security vulnerability in SSL.
- Corrects a potential security vulnerability when using the Secure Sockets Layer (SSL). The potential vulnerability may be remotely exploitable, resulting in a denial of service (DOS).

## Patch 27115.00

*OSFSSOW2K540*

- Fixes memory leaks in the libgssldap library used by ldapcd. The leaks caused intermittent SSO SIA authentication failures.
- Corrects an "address already in use" problem with klogin and kshell.
- Fixes a problem that occurs when running a GSSAPI application, where instead of returning error-specific strings, generic error-strings are returned.



---

## 3 TruCluster Server Patches

This chapter provides information about the defect fixes included in Version 5.1B-6 for the TruCluster Server software. This release contains only defect fixes. There are no new features or enhancements added in this release of TruCluster Server software. It is organized as follows:

- “New Release Notes” lists release notes that are specific to the TruCluster Server software patches in this kit and TruCluster Server issues in general.
- “Prior Release Notes” section lists release notes listed from the initial Version 5.1B release through Version 5.1B-5.
- “Summary of TruCluster Server Software Patches” provides brief descriptions of TruCluster Server patches that are new to this kit and “Patches Delivered in Previous Kits” describes patches that were first delivered in previous Version 5.1B kits and are included in this kit.

### New Release Notes

This release contains only defect fixes. There are no new features or enhancements added in this release. For important information about installing or removing previous versions, see “Cluster-Specific Installation and Removal Release Notes”.

### Prior Release Notes

The following sections describe some of the key features and enhancements that were first delivered in previous patch kits.

### New Flag Option to Turn OFF All Existing Flags for Services

A new cluster alias option `none` can be added for a service in `/etc/clua-services` files. This option allows you to unset an option without needing to specify a different option. The most common use of `none` is to unset the option of a service that has only one option set.

### Select Option to Check Tagged Files

During the preinstall stage of a rolling upgrade, you have the option of checking tagged files. You should override the default setting and select the check tag option. The reason for selecting this option is described in “Check for Tagged Files if Messages Are Displayed”.

### Check for Tagged Files if Messages Are Displayed

When installing this patch kit during a rolling upgrade, you may see the following error and warning messages during the setup stage:

```
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/laddebug.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.
```

If you see these messages during the setup stage, you should verify that the tagged files were properly created when you execute the preinstall stage.

In cases where the tagged files are not created, you can repeat the setup stage.

## Noncritical Errors

During a rolling upgrade to install this patch kit, you may encounter the following noncritical situations:

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

## 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.

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

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

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

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

## Undo Stages in Correct Order

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

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

## 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 the [Patch Kit Installation Instructions](#) and the `installupdate(8)` and `clu_upgrade(8)` reference pages.

## 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.

## Relocating AutoFS During a Rolling Upgrade on a Cluster

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

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

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

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

```
caa_stat -t
Name Type Target State Host

autofs application ONLINE ONLINE rye
cluster_lockd application ONLINE ONLINE rye
clustercron application ONLINE ONLINE swiss
dhcp application ONLINE ONLINE swiss
named application ONLINE ONLINE rye
```

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

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

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

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

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

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

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

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

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

## Messages Displayed During Rolling Upgrade Can Be Ignored

You can ignore the following messages if you see them displayed during a rolling upgrade:

- kill:1048674: no such process

This message may be displayed after the roll stage. For example:

```
clu_upgrade roll
```

```
This is the cluster upgrade program.
The 'roll' stage has completed successfully. This
member must be rebooted in order to run with the newly
installed software.
Do you want to reboot this member at this time? []:y
You indicated that you want to reboot this member at this time.
Is that correct? [yes]:
```

```
The 'roll' stage of the upgrade has completed successfully.
kill: 1048674: no such process
```

```
#
```

- rmdir: /var/.clu\_upgrade: File exists

This message may be displayed after the clean stage. For example:

```
clu_upgrade clean
```

```
This is the cluster upgrade program.
```

```
You have indicated that you want to perform the 'clean' stage
of the upgrade. Do you want to
continue to upgrade the cluster? [yes]:
```

```
Deleting tagged files.
```

```
.....
.....
.....
.....
.....Removing back-up and kit files
```

```
rmdir: /var/.clu_upgrade: File exists
```

```
The 'clean' stage of the upgrade has completed successfully.
```

```
#
```

## Error on Cluster Creation

When you attempt to create a cluster after having deleted patches, you may see the following error messages:

```
*** Error ***
```

```
This system has only Tru64 UNIX patches installed.
Please install the latest TruCluster Server patches on your system.
You can obtain the most recent patch kit from:
http://www.support.compaq.com/patches/
```

```
*** Error ***
```

```
The system is not configured properly for cluster creation.
```

Please fix the previously reported problems, and then rerun the 'clu\_create' command.

If you see these messages, enter the following command:

```
ls -t1r /usr/.smdb./*PAT*.sts
```

If this command returns a file with 000000 in its name, you will have to run the `clu_create` command with the `-f` option to force the creation of your cluster. The problem is caused by the cluster software misinterpreting the existence of some patches and will be corrected in a future patch kit.

If the command does not return a file with 000000 in its name, you will need to contact HP support to determine the cause of the problem.

## 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, enter the following commands:

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

## Login Failure Possible with C2 Security Enabled

Login failures may occur as a result of a rolling upgrade on systems with Enhanced Security (C2) enabled. The failures may be exhibited in two ways:

- With the following error message:  
Can't rewrite protected password entry for user
- With the following set of error messages:  
login: Ignoring log file: /var/tcb/files/dblogs/log.00001: magic number 0, not 8  
  
login: log\_get: read: I/O error  
Can't rewrite protected password entry for user

The problem may occur after the initial reboot of the lead cluster member or after the rolling upgrade is completed and the `clu_upgrade switch` procedure has been run. The following sections describe the steps you can take to prevent the problem or correct it after it occurs.

### Preventing the problem

You can prevent this problem by performing the following steps before beginning the rolling upgrade:

1. Disable the `prpasswd` daemon from running on the cluster:

```
rcmgr -c set PRPASSWDD_ARGS \
"rcmgr get PRPASSWDD_ARGS` -disable"
```
2. Stop the `prpasswd` daemon on every node in the cluster:

```
/sbin/init.d/prpasswd stop
```
3. Perform the rolling upgrade procedure through the `clu_upgrade switch` step and reboot all the cluster members.
4. Perform one of the following actions:

- If PRPASSWDD\_ARGS did not exist before this upgrade (that is, if `rcmgr get PRPASSWDD_ARGS` at this point shows only `-disable`), then delete PRPASSWDD\_ARGS:
 

```
rcmgr -c delete PRPASSWDD_ARGS
```
  - If PRPASSWDD\_ARGS existed before this upgrade, then reset PRPASSWDD\_ARGS to the original string:
 

```
rcmgr -c set PRPASSWDD_ARGS \
" `rcmgr get PRPASSWDD_ARGS | sed 's/ -disable//'`"
```
5. Check that PRPASSWDD\_ARGS is now set to what you expect:
 

```
rcmgr get PRPASSWDD_ARGS
```
  6. Start the prpasswd daemon on every node in the cluster:
 

```
/sbin/init.d/prpasswd start
```
  7. Complete the rolling upgrade.

## Correcting the problem

If you have already encountered the problem, perform the following steps to clear it:

1. Restart the prpasswd daemon on every node in the cluster:
 

```
/sbin/init.d/prpasswd restart
```
2. Reboot the lead cluster member.
3. Check to see if the problem has been resolved. If it has been resolved, you are finished. If you still see the problem, continue to step 4.
4. Try to force a change to the auth database by performing the following steps:
  - a. Use `edauth` to add a harmless field to an account, the exact commands depend on your editor. For example, pick an account that does not have a vacation set and add `u_vacation_end`:
 

```
edauth
s/:u_lock@:/u_vacation_end#0:u_lock@:/
w
q
```
  - b. Check to see that the `u_vacation_end#0` field was added to the account:
 

```
edauth -g
```
  - c. Use `edauth` to remove the `u_vacation_end#0` field from the account.

If the `edauth` commands fail, do not stop. Continue with the following instructions.
5. Check to see if the problem has been resolved. If it has been resolved, you are finished. If you still see the problem, observe the following warning and continue to step 6.



### Warning!

Continue with the following steps only if the following conditions are met:

- You encountered the described problem while doing a rolling upgrade of a cluster running Enhanced Security.
  - You performed all previous steps.
  - All user authentications (logins) still fail.
- 

6. Disable logins on the cluster by creating the file `/etc/nologin`:
 

```
touch /etc/nologin
```
7. Disable the prpasswd daemon from running on the cluster:
 

```
rcmgr -c set PRPASSWDD_ARGS \
" `rcmgr get PRPASSWDD_ARGS` -disable"
```

8. Stop the `prpasswd` daemon on every node in the cluster:
 

```
/sbin/init.d/prpasswd stop
```
9. Force a checkpoint of `db_checkpoint`, using the `db_checkpoint` command with the `-1` (number 1) option :
 

```
/usr/tcb/bin/db_checkpoint -1 -h /var/tcb/files
```

 Continue with the instructions even if this command fails.
10. Delete the files in the `dblogs` directory:
 

```
rm -f /var/tcb/files/dblogs/*
```
11. Force a change to the auth database, as follows:
  - Use the `edauth` command to add a harmless field to an account, the exact commands depend on your editor. For example, pick an account that does not have a vacation set and enter the following:
 

```
edauth
s/:u_lock@:/u_vacation_end#0:u_lock@:/
w
q
```
  - Check to see that the `u_vacation_end#0` field was added to the account:
 

```
edauth -g
```
  - Use the `edauth` command to remove the `u_vacation_end#0` field from the account.



### Warning!

If the `edauth` command fails, do not proceed further. Contact HP support.

12. If the `edauth` command was successful, perform one of the following actions:
  - If `PRPASSWDD_ARGS` did not exist before this upgrade (that is, if `rcmgr get PRPASSWDD_ARGS` at this point shows only `-disable`), then delete `PRPASSWDD_ARGS`:
 

```
rcmgr -c delete PRPASSWDD_ARGS
```
  - If `PRPASSWDD_ARGS` existed before this upgrade, then reset `PRPASSWDD_ARGS` to the original string:
 

```
rcmgr -c set PRPASSWDD_ARGS \
 "`rcmgr get PRPASSWDD_ARGS | sed 's/ -disable//'\`"
```
13. Check that `PRPASSWDD_ARGS` is now set to what you expect:
 

```
rcmgr get PRPASSWDD_ARGS
```
14. Start the `prpasswd` daemon on every node in the cluster:
 

```
/sbin/init.d/prpasswd start
```
15. Re-enable logins on the cluster by deleting the file `/etc/nologin`:
 

```
rm /etc/nologin
```
16. Check to see if the problem has been resolved. If it has not, contact HP support.

## File System Unmount Recommended if Message Is Displayed

Under certain error conditions, the following message may be seen during a relocation or failover, or during the boot of a member:

```
WARNING: Unable to failover /mnt: pfs and cfs fsids differ
```

The result is that the fileset in question is now unserved in the cluster. For example:

```
cfsmgr /mnt
Domain or filesystem name = /mnt
Server Status : Not Served
```

If this occurs, we recommend that you immediately do the following:

1. Use the following command to unmount the filesystem:

```
cfsmgr -u -p [mountpoint]
```

2. If other mounted filesets exist in the same domain, unmount them (they should also be in the "Not Served" state):

```
cfsmgr -u -d [domain]
```

For steps on checking an AdvFS domain, see the AdvFS Administration Guide, Section 6.3.1, steps 3-7.

3. Run diagnostics on the domain prior to remounting its file systems.

To verify the domain, you can use the AdvFS `verify` utility or the `fixfdmn` utility. If using `fixfdmn`, we recommend first running it with the `-n` option to see what errors are found prior to allowing `fixfdmn` to fix them.

Once you have successfully verified the domain, remounting the domain's file systems in the cluster should succeed.

If the domain cannot be immediately verified, we recommend that you do not remount the original fileset until this can be done.



#### Note:

In rare cases, the warning message will be accompanied by a system panic. This will occur if CFS error handling is unable to successfully unmount the underlying physical file system. If this occurs, the console will direct you to use `cfsmgr` to unmount the domain on one of the remaining nodes prior to rebooting the member.

This action will prevent the rebooted member from attempting to failover-mount the file system and will minimize access to the domain. Prior to remounting the file system, it is advisable that the domain be sanity-checked using the steps given above.

---

## Tunable Attribute May Help Performance Problem

The tunable attribute `cfs_clone_noccr`, included in this patch kit, may correct a problem in which cluster fileset writes that occur simultaneously with reads of the fileset clone on a cluster client (for example, during a backup) may result in performance degradation. This occurs most often when the clone file being read consists of many thousands of extents (for example, 20,000 or more).

If a degradation during cluster clone reads is noticeable (for example, the clone read appears to be hanging and requires a long time to complete), set the value of `cfs_clone_noccr` to 1 on the server of the given fileset. This sysconfig tunable attribute is set to 0 by default and should be changed only when the degradation is noticeable.

Note that all filesets with clones that are served by the node on which the attribute is set will also see this change. It may be advisable (though not required) to have those filesets whose clone files have fewer extents be served by a different node during the time the tunable attribute is set.

## AlphaServer ES47 or AlphaServer GS1280 Hangs When Added to Cluster

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

Use the following instructions to extract and copy the V5.1B cluster `genvmunix` from your original Tru64 UNIX kit to your AlphaServer ES47 or AlphaServer GS1280 system. In these instructions, the AlphaServer ES47 or AlphaServer GS1280 is designated as member 5. Substitute the appropriate member number for your cluster.



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

```
mount -r /dev/disk/cdrom0c /mnt
```
3. Mount the boot disk of the AlphaServer ES47 or AlphaServer GS1280 on its specific mount point; for example:  

```
mount root5_domain#root /cluster/members/member5/boot_partition
```
4. Extract the original clu\_genvmunix from the CD-ROM and copy it to the boot disk of the AlphaServer ES47 or AlphaServer GS1280 member.  

```
zcat < TCRBASE540 | (cd /cluster/admin/tmp; \
tar -xf - ./usr/opt/TruCluster/clu_genvmunix)
#cp /cluster/admin/tmp/usr/opt/TruCluster/clu_genvmunix \
/cluster/members/member5/boot_partition/genvmunix
rm /cluster/admin/tmp/usr/opt/TruCluster/clu_genvmunix
```
5. Unmount the CD-ROM and the boot disk:  

```
umount /mnt
umount /cluster/members/member5/boot_partition
```
6. Reboot the AlphaServer ES47 or AlphaServer GS1280.

## 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`.

## Data Protector Issues and Restrictions

The following sections describe issues and restrictions for Version 5.1 of the HP OpenView Storage Data Protector backup and recovery product when configuring it on a Tru64 UNIX cluster.

### Possible Error Backing Up Cluster Mount Points

When backing up cluster mount points using the cluster alias as the client name, you may encounter an error in which the directory is reported as a mount point to a different file system and is backed up as an empty directory.

To correct this problem, create TruCluster Server clients as follows:

- Create a client for each host name node in the cluster.
- Create another client using the cluster alias name, selecting it as a virtual host.

You can then create backups using the alias as the client name.

You may also need to define your mount points to back up using the manual add function of the Add Backup wizard. Under some circumstances, backups that are created using the default device discovery encounter the "backed up as an empty directory" problem.

## Configuring Data Protector for Oracle Integration

When Configuring Data Protector for Oracle integration, `libobk.so` should be linked with `/usr/omni/lib/libob2oracle8_64bit.so`.

The *Data Protector UNIX Integration Guide* incorrectly states that it should be linked with `/usr/omni/lib/libob2oracle8_64.so`.

## Set `ippport_userreserved` Attribute on Large Systems

Larger systems can encounter portmapper problems in a local area network (LAN) cluster if the value of the `ippport_userreserved` attribute has not been tuned. The recommended value is 65535 and should be the same for all cluster members. Set the value before adding the first member.

If this value is not set for a LAN cluster with larger machines, the machines may run out of ports for interconnect services. For more information, see the manual *Tuning Tru64 UNIX for Internet Servers*.

## Summary of TruCluster Server Software Patches

The following sections provide brief descriptions of the changes delivered in this patch kit and in previous Version 5.1B patch kits for the TruCluster Server software products.

Each patch provides fixes to subsets of the operating system. Subset names (listed in italic font in the following list) consist of three parts; for example, for subset *TCRBASE540*, the *TCR* indicates that the subset is part of the TruCluster Server product, the *BASE* indicates a category, and the *540* indicates that the subset belongs to the Version 5.1B operating system.

### New Patches

The patch summaries in this section describe defect fixes new to this release. There are no new features or enhancements added in this release.

#### Patch 29001

##### *TCRBASE540*

- Fixes system panics while unconfiguring `cnx` subsystem.
- Corrects CFS to support volume names of size up to the `pathnames` size.
- Fixes a KMF during a CFS domain relocation.
- Fixes such that UFS deadlocks are avoided during file/directory rename operations.
- Fixes a problem in `cnx tunable qdisk_open_delay`.

#### Patch 29002

##### *TCRMAN540*

- Updates `sys_attrs_cnx(5)` reference page.

### Patches Delivered in Previous Kits

The following TruCluster Server patches were delivered in previous Version 5.1B patch kits. These patches will be installed on your system if you did not install the previous kit.

#### PATCH 28001.00

##### *TCRBASE540*

- Fixes a problem in which a CFS client read operation returns the wrong data due to stale metadata associated with the file frag.
- Adds a check to prevent the caller from binding to a cluster alias address that the node has not joined.

- Fixes an infinite loop under certain circumstances in `cms_do_mount_rpc()`.
- Added option to unset all flags for a service in `/etc/clua_services`.
- Corrects a reference count issue in the KGS subsystem.
- Fixes node panic with `ics_unable_to_make_progress: netisrs stalled`, though `net_isr` thread was not actually stalled.
- Provides a fix for a domain panic caused by hung IOs on a busy or faulty disk drive. The panic can happen after all but one path to the disk drive being disabled then re-enabled.
- Corrects a problem where a 'local open' on a previously opened tape drive results in an erroneous "no such device" message.
- Provides a fix for a cluster boot-time hang, caused by a fault quorum disk.
- Fixes multiple issues with RDG(Reliable DataGram) component in a LAN cluster.
- Fixes an issue with CFS failover subsystem where, under certain domain configurations failover process may hang.
- Fixes a problem with `fuser(8)` where usage of the `-a` option leaves the filesystem incapable of unmounting even if no files or directories on the filesystem are in use.
- Fixes a problem where, under certain circumstances, a close on socket of type `AF_UNIX` may result in a system panic.
- Provides enhancements to the DRD trace framework.
- Optimize the performance of `ics0` interface in a LAN cluster.
- Fixes an issue with aliasd routing in a cluster.
- Avoids panic due to bad quorum disk during boot process.
- Fixes an issue wherein Internode Communication Subsystem panics when it receives messages for an unknown service.
- Updates `volstat` utility and kernel to report cluster-wide LSM statistics.
- Add support in cluster alias to handle socket unlisten.

## PATCH 28002.00

### TCRMAN540

- Provides the latest reference pages for `sys_attrs_cfs(5)`, `sys_attrs_clubase(5)`, and `sys_attrs_rdg(5)`.
- Updates `clu_alias.config(4)` and `exports.alias(4)` reference pages.
- Updates `sys_attrs_icsnet(5)` reference page to reflect `icsnet_mtu` attribute.
- Updates the following reference pages: `clua_services(4)`, `cfsd.conf(4)`, `sys_attrs_ics_ll_tcp(5)`
- Updates the following reference pages: `imcs(1)`, `dln_rd_collect(3)`, `dln_rd_validate(3)`, `imc_rderrcnt(3)`, `sys_attrs_cms(5)`, `sys_attrs_drd(5)`, `sys_attrs_icsnet(5)`

## Patch 27001.00

### TCRBASE540

- Eliminates numerous panics and hung devices by fixing `drd` so it no longer accesses a device that has a deletion pending or in progress.
- Fixes an RM simple lock timeout issue that may occur in noisy Memory channel rails.
- Enhances the error message generated when the `clu_bdmgr` command cannot access a member boot disk.
- Fixes a configuration issue found in non-CAM devices and CD\_ROM devices.
- Fixes the cause of potential cluster hangs during some Memory Channel hardware failures that result in an MC rail failover.
- Fixes the CFS AIO write error path so the I/O completion steps are not repeated.
- Fixes a flaw in CFS file locking code that causes a "vrele: bad ref count" panic.
- Fixes the cause of an assertion failure in `cfs_vnops.c`.

- Corrects a problem in which the simultaneous booting of multiple nodes results in a panic due to an unknown node in a remote member node list.
- Corrects a problem in a Memory Channel cluster in which a panic occurs in a booted member when a booting member goes down because of panic/halt/shutdown.
- Fixes a problem in which a thread enters dio code while an extent map is being refreshed.
- Fixes a problem of v\_numoutput not decremented for aio dio error paths.
- Removes the cause of a panic that may occur in CFS at boot time if a remote node goes down.
- Corrects several ICS signal-forwarding issues.
- Fixes a race between the close system call for a block device file and the recovery process for the file system.
- Clarifies a usage message seen with the cfsstat command.
- Corrects a problem in clu\_mibs daemon that can cause various eSNMP sub agents, such as pmgrd and os\_mibs, to terminate.
- Fixes a problem to prevent the relocation of UFS read/write file system to the original node.
- Provides new option to the mountd daemon to specify a port number for mountd to bind to.
- Corrects a problem in which a DRD event thread may run infinitely while responding for bid server transaction.
- Fixes and AdvFS domain panic caused by cfsd.
- Corrects a problem in CAA in which a resource does not fail over when two resources have the same values for the FAILOVER\_DELAY and REQUIRED\_RESOURCES attributes.
- Fixes a hang during cluster bootup caused by early reservation conflicts.
- Provides enhancements to the caa\_relocate command.
- Provides a new command, clu\_ping, to determine the status of the interconnects in a stretched cluster environment.
- Improves CFS client writing to do the following:
  - reduce the logging of ERROR 69 for user disk space quota exhausted.
  - support partial write success.
  - increase the interconnect transfer size for multi-page synchronous writes.
  - prevent read ahead past the end of a file.
- Helps ensure more accurate block reservation accounting in CFS.
- Addresses an issue seen on Tru64 UNIX LAN clusters, whereby a booting node may panic with "lock\_wait" while spawning threads for cluster interconnect channels.
- Provides a solution to display a warning message if deleting a particular cluster member would cease NTP services for the rest of the cluster.
- Improves the routing fail-over mechanism when one or more network interfaces on any cluster member fails.
- Fixes a "kernel memory fault" panic in cfs\_fo\_failover\_done().
- Fixes a problem wherein the DRD subsystem may cause a system panic when strategy routines are called from a light weight context (LWC).
- Fixes display errors in the cfsstat command when using the icschanbps option.
- Fixes display errors in cfsstat command when using the icschanbps option.
- Fixes a deadlock issue between cluster nodes because of cfs\_async\_io\_thread running on them.
- Corrects an erroneous error message displayed by drdmgr.
- Fixes a cnx\_qdisk\_thread hang problem.
- Fixes a memory leak in CFS.
- Fixes disk I/O hang in DRD.
- Fixes a hang with disklabel that occurs if a local open fails for the same disk simultaneously.

- Fixes incorrect CFS token structure warnings.
- Prevents file inconsistency due to a race between lookup and remove.
- Provides a new cluster-specific link aggregation distribution algorithm when using LAG in a LAN cluster.
- Fixes a simple lock timeout panic issue in kch and a possible hang at boot time
- Prevents an AIO DirectIO to return invalid data while reading a fragged file.
- Fixes a cluster hang issue during cluster boot-up, when local disk open operations fail while disklabel is in progress.
- Fixes an error in the DRD subsystem wherein un-initialized disk attributes can cause a system panic.
- Fixes KMF in `rdg_get_completion()` routine.
- Fixes a problem in which a cluster alias subsystem tries to free the mbuf that is already freed by ICS subsystem.
- Corrects reference counting issues within the DRD subsystem that can prevent the deletion of hwids.
- Adds a new option, `custom_gated`, to `cluamgr` and `aliasd`.
- Fixes a deadlock that can happen during failover of global root and var file systems when `vfast` is enabled on them.
- Fixes resource leaks seen after a locked device file is revoked.
- Fixes system panics seen on relocating file systems with locked revoked devices.
- Fixes a problem with CAA placement policy when host names in "HOSTING MEMBERS" are in uppercase letters.
- Corrects a problem in which CAA is incorrectly showing the status of network resources on a halted member.
- Fixes a problem in `cfs` block reservation code where `cfs` attempted to release a lock more than once.
- Introduces a code tracing capability of the `aliasd` and `aliasd_niff` daemons to improve troubleshooting.
- Prevents a race that can occur during the planned relocation of a file system.
- Improves the reliability of the DRD subsystem when faced with tape devices and tape device failures.
- Introduces a mechanism to improve reliability for synchronizing cluster alias ID sets among cluster members.
- Fixes the cause of the following CNX panic in cluster reconfiguration:  
`cnx_change_cluster_tx_state: illegal transaction state`
- Fixes an ICS panic issue that occurs early in the boot process.
- Fixes a problem that causes the cluster alias manager `SUITlet` to falsely interpret any cluster alias with `virtual={t|f}` configured as a virtual alias regardless of its actual setting.
- Corrects problems in which `SysMan drdmgr` dumps `tcl` stack when a user tries to manage devices or file systems of a cluster node that is down.
- Corrects an issue to allow the Device Request Dispatcher, DRD, to retry to get disk attributes when `EINPROGRESS` is returned from the disk driver.
- Address issues with "address already in use" messages from `klogin` and `kshell`.
- Corrects a potential security vulnerability in CAA.
- Fixes a kernel memory fault.
- Corrects a problem in which the MC-API call `imc_ckerrcnt_mr()` incorrectly returns an error status, although the functions error count parameter is not increasing.
- Preserves the error code from an asynchronous write error on a CFS client and returns the error from the `close()` system call.

- 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.
- Corrects a problem to allow the use of 255 in the LAN Interconnect IP address.
- Fixes a CFS client panic during a file system read operation where the server goes down and the client itself becomes the server and attempts to release the direct I/O token that had already been released.
- Fixes a forced unmount of nonfailoverable file system (that is, NFS and AutoFS) panic in the case that the initiator is down.
- Enables a cluster to boot even if the cluster root domain devices are private to different cluster members. Although this is not a recommended configuration, it should not result in an unbootable cluster. Currently, this is with respect to cluster root domains not under LSM control.
- Corrects a potential data inconsistency caused by a problem in the CFS block reservation code, which calculates incorrectly the amount of space requested and used by direct I/O writes.
- Resolves a kernel memory fault in `m_copym`.
- Fixes a problem with the `-b` option of `caa_report`.
- Fixes a problem with `caa_stop -f` by allowing the administrator to reset a resource state from UNKNOWN to OFFLINE even if the hosting member is down.
- Corrects a potential data inconsistency that may occur when a domain is nearly full. Client write requests shipped synchronously to the server will no longer have subsets of pages written asynchronously due to a race with virtual memory.
- Improves the scaling of IP reassembly code on large SMP machines. NFS servers are especially susceptible when a large number of clients attempt to write at the same time.
- Helps to close a race where synchronous writes may obtain disk allocations that were promised to cached client writes.
- Fixes a problem in which CAA might prevent alias based services from properly functioning by binding to one the cluster alias reserved ports.
- Corrects a problem in a Memory Channel cluster where rebooting a node without performing a hardware reset can crash other members with a `RM_AUDIT_ACK_BLOCK` panic.
- Fixes a problem in the Memory Channel driver.
- Improves the responsiveness of EINPROGRESS handling during the issuing of I/O barriers by removing a possible infinite loop scenario that could occur due to the deletion of a storage device.
- Fixes a problem that causes a panic with the message "CNX MGR: Invalid configuration for cluster seq disk" during simultaneous booting of cluster nodes.
- Fixes the panic "CNX MGR: Invalid configuration for cluster seq disk" that occurs during the simultaneous booting of cluster nodes.
- Fixes a possible race condition between a SCSI reservation conflict and an I/O drain that can result in a hang.
- Alleviates a condition in which a cluster member takes an extremely long time to boot when using LSM.
- Fixes a problem that `caa_relocate` AutoFS does not kill the `autofs` daemon.
- Allows rewrites when the domain is close to out of space.
- Ensures correct processing in the `close()` system call.
- Provides a CAA action script that can be used by a NIS Slave running to help assign a crontab entry to update NIS maps.
- Fixes a problem in which a cluster member leaves the cluster alias yet continues to respond to it.
- Corrects a problem that causes applications (including `cluamgr`) to get a dummy cluster alias reported from the `cluaioc_get_nextalias()` call. The IP address for this alias is 0.0.0.0.

- Fixes a problem in which aliasd creates multiple similar virtual subnet static routes in the `gated.conf.memberX`, thereby causing `gated` to fail to load.
- Fixes issues associated with the initialization of the Memory Channel driver.
- Provides a function to query the status of `aliasd`.
- Fixes an IPv6 bind problem in a cluster environment.
- Fixes multiple disable or enable problems with `cluamgr`.
- Fixes a `tok_wait` hang problem on Sierra Clusters.
- Adds the ability to change the default interconnect interface name.
- Corrects several problems in the cluster install and upgrade utilities.
- Fixes a problem in which an RDG (Reliable DataGram) kernel thread can starve other timeshare threads on a uniprocessor cluster member. In particular, system services such as networking threads can be affected.
- Fixes minor issues with `cfsstat` command-line options and return values.
- Prevents panics seen with cluster server-only (for example, MFS) mounts.
- Fixes a condition that causes the panic `pg_nwriters` going negative when `ubc_page_release()` is called from `cfs_getpage()`.
- Corrects a problem in the RDG component in which multiple Oracle instances are unable to be properly configured when using RDG over a LAN rather than Memory Channel.
- Provides a sticky connection feature for a cluster alias.
- Updates `sysconfig` to use the cluster interconnect, allowing for a greater SSI collaboration. This will help with changing variables on hung systems, single user systems, and normal running systems.
- Improves device error processing in `drd`.
- Corrects a boot hang problem seen on large-scale Sierra Cluster configurations caused by a missed wake up in the kernel group services code.
- Alters the behavior of the cluster NFS client with TCP mounts so that when a remote server is down, the cluster NFS client will use nonreserved ports to see if the remote server is up.
- Introduces a new CFS tunable attribute that may benefit the performance of client reads of clone files under certain circumstances.
- Addresses an assertion caused by a bad user pointer passed to the kernel via `sys_call`.
- Corrects a condition that results in excessive context switching and CPU load due to a heavy use of the cluster alias on large SMP and NUMA machines .
- Enhances `/sbin/advfs/tag2name` to print out the name of the associated directory, given the tag of an index file.
- Increases performance scalability and extends the reliability of the Internode Communications Subsystem in a cluster configured with Memory Channel as the cluster interconnect.
- Improves detection of possible race conditions during CFS recovery.
- Adds a cluster panic facility to the kernel.
- Addresses the following:
  - An issue in which new ICS server daemons and handles are created one at a time each time the low water mark for each is reached, thereby causing a nanny daemon to be called more frequently than it needs to.
  - An issue in which no mechanism exists for the user to adjust the high and low water marks for ICS free handles, which can result in poor performance during rapidly increasing loads.
- Fixes a problem in which cluster alias connections are not distributed among cluster members according to the defined selection weight.
- Fixes a memory leak in the cluster alias subsystem.
- Fixes an issue with ICS (Internode Communication Services) on a NUMA-based system in a cluster.

- Fixes a problem in the cluster kernel in which a cluster member panics while doing remote I/O over the interconnect.
- Fixes a hang that occurs when multiple nodes are shutting down simultaneously; fixes a Cluster File System panic that occurs when using raw Asynchronous I/O; and provides additional code to assist in problem diagnosis.
- Corrects a problem in which a panic displaying the message “error CNX MGR: cnx\_comm\_error: invalid node state” occurs on a LAN cluster running under load when other members are rebooting.
- Addresses an error in which `caa_register -u` produces with no balance data.
- Addresses a resource inaccessibility issue that can occur if the hosting member crashes during a remote `caa_stop` operation.
- Updates the attributes on a directory when files are removed by a cluster node that is not the file system server.
- Fixes a problem associated with non-SCSI storage.
- Corrects a potential security vulnerability in the cluster interconnect security configuration that may result in a denial of service (DoS) on systems running TruCluster Server software.
- Causes UDP datagrams that do not come from the correct port to be discarded.
- Addresses a node hang that occurs during the testing of Memory Channel cable pulls. A cluster member may hang when a Memory Channel cable is pulled, the node is taken down, the cable is plugged back in, and the node is rebooted.
- Fixes a cluster deadlock that may occur during a failover and recovery when direct I/O is in use.
- Fixes a race condition in the Device Request Dispatcher.
- Corrects a condition that can cause excessive `FIDS_LOCK` contention when a large number of files are using system-based file locking.
- Fixes a problem with `cfds` core dumping shortly after startup if it is enabled or shortly after enabling it. The problem fixed by this patch is only seen after applying a recent `dsfmgr` patch.
- Corrects diagnostic code that could cause a panic during a kernel boot.
- 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.
- Prevents a panic when an AutoFS file system is auto-unmounted.
- Corrects the cause of a cluster member panic with kernel memory fault when running `nmap` or `nessus` targeting at the cluster alias.
- Resolves a problem in which the `caa_register` command allows a CAA resource to be registered even when its profile contains an unknown attribute. This fix prevents the `caa_register` command from registering a resource with an unknown attribute and will cause it to return an error message that includes the unknown attribute information.
- Fixes a condition in which uptimes greater than 100 percent are reported for resources by `caa_report`.
- Fixes a problem in which resources that never started have an ending timestamp.
- Fixes a problem in which CAA dumps core when trying to deal with cluster member ID 63.
- Fixes an problem where access to the quorum disk can be lost if the quorum disk is on a parallel SCSI bus and multiple bus resets are encountered.
- Relieves pressure on the CMS global DLM lock by allowing AutoFS auto unmounts to back off.
- Fixes `cfsmgr` to properly return a failure status when a relocation request has failed.
- Fixes a race condition where stale name cache entries allow file access after file unlink.
- Corrects a problem in which `cfds` will terminate prematurely and core dump when a node leaves the cluster very shortly after joining the cluster.
- Fixes a timing window during asynchronous reads on a CFS client.



- Fixes a panic that may occur during an unmount.
- Corrects several problems with various installation commands and utilities.
- Fixes a memory leak in the `clu_get_info` interface.
- Enhances cluster file system performance when using file locks to coordinate file access.
- Causes the correct error message for `freezefs -q` to be displayed on a non-AdvFS file system.
- Fixes a problem in one of the shipped rc scripts whereby Oracle fails during startup on a clustered system.
- Addresses a panic that occurs on a booting node.
- Fixes a coding error, a memory leak, and a deinitialization problem in the cluster interconnect networking layer.
- Fixes a problem in the Device Request Dispatcher.
- Provides `clu_upgrade` enhancements.
- Increases performance by reducing the lock miss rate in the `ics_mct_llnode_info_lock`.
- Addresses the panic "Assert Failed: (cp-c\_flags & CDIRECTIO) = 0" in the cluster file system.
- Corrects a problem where a CFS lookup for a mount could leave stale state behind that could adversely affect subsequent NFS operations.
- Fixes an internal problem in the kernel's AdvFS, UFS, and NFS file systems where extended attributes with extremely long names, greater than 247 characters, could not be set on files. The new limit is 254 + a null string terminator.
- Corrects problems with LSM disks and the cluster quorum tools. When a member having LSM disks local to it is down, the quorum tools fail to update quorum. This causes other cluster commands to fail.
- Corrects a problem in which mounting on a directory in a clone fileset fails with the message "Device Busy."
- 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 which CAAD might dump core due to a race condition when multiple events to which it subscribes arrive simultaneously.
- Improves the fragment gathering mechanism to boost performance.
- Fixes panic problem when attempting to unload `clua.mod`.
- Fixes a condition that causes a boot up panic when `ippport_userreserved` is 1000 or less.
- 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.
- Optimizes cluster file system lock recovery, potentially speeding up the time required to failover a file system to a new server.
- Corrects a condition in which superfluous "rm\_event, index too big" messages may appear on system consoles.
- Addresses a panic that may occur when a node is joining the cluster. A node recognizing the joining node panics while it is trying to establish a preboot channel connection with the peer node, causing the following message to be displayed on the console or in `/var/adm/messages`:  
panic (cpu x): ics\_mct: rx conn 3
- Corrects the LSM partition types in the CNX partition of boot disk for the `clu_partmgr` utility.
- Modifies the `aliasd` daemon to include interface aliases when determining whether or not an interface is appropriate for use as the ARP address for a cluster alias when selecting the proxy ARP master.
- Fixes the potential of multiple `assert_wait` and `timeout` panics due to kernel EVM threads not properly preempting.

- Fixes a problem in the Memory Channel driver.
- Corrects a condition that occurs during a rolling upgrade in which the `clu_ifaccess` script removes the tag file for `/etc/ifaccess` and sends out a warning message.
- Forces a reboot to resolve communications problems in a two node cluster rather than hang.
- Corrects lock acquires after `mpsleap`.
- Causes a rebuild delay remainder to be minimally second.
- Allows the cluster to provide new functions to the `dupatch` command before a member is rolled, and also provides a mechanism for backing out the added functions.
- Addresses a memory leak in the Memory Channel transport layer.
- Fixes a problem in which a system may panic with a kernel memory fault when a device that is being opened by one program is being deleted with the `hwmgr` utility.
- Fixes a condition that causes a panic when a valid NFS packet with corrupted embedded length field is received.
- Fixes a condition that causes an unnecessary panic due to request connection deregistration with an invalid IP address.
- Provides performance improvement for CFS filesets mounted with the `server_only` option. A log sync for create transactions is not needed for such filesets.
- Fixes a problem with single physical rail Memory Channel configurations and cleans up stale data left on an off-line physical rail by the Memory Channel driver.
- Fixes a rare cluster hang caused by dead locks that occurred between the CFS client and server during multiple write operations.
- Fixes multiple problems seen with the TruCluster RDG component, including panics of the following types "rdg: unwiring", "vl\_unwire: page is not wired", and "KMF: from \_otsmove."
- Allows users to add new members and create a cluster with different netmasks.
- Removes member0-specific installation files on an undo install, which could prevent the reinstallation of the patch.
- Allows users to continue forward when they add a member to a one-node cluster during a rolling upgrade or rolling patch.
- Enables CAA to start up and fail over system services before any of the user services.
- Fixes an unaligned kernel access in the cluster I/O stack.
- Addresses a potential hang in the NFS server that occurs when file systems are being relocated in a cluster.
- Provides the ability to lower the `cluster_rebuild_delay`.
- Fixes the long delay during an NFS connection failover when servicing cluster member dies.
- Fixes a panic in `clua.mod` that is caused by receiving a `delete-cnx-request` from a member when that `cnx` is in the UNREGISTER state.
- Fixes a reconnection problem when an interface comes down and then goes up.
- Fixes a panic problem in `clua.mod` that occurs when `max_aliasid` is increased and aliases are added.
- Fixes a situation that causes a core dump in `aliasd` when all interfaces are removed on a cluster member that is set up with at least one cluster alias that was added with `virtual=t` and without a subnet.
- Fixes a problem when disabling and re-enabling cluster alias source route on a given interface.
- Fixes a problem where `clua.mod` does not handle TCP RST messages appropriately.
- Fixes a problem of restoring static routes when an interface revives.
- Corrects a problem in which a rolling upgrade stops advancing when adding a cluster member to a one-node cluster.
- Fixes an initialization issue with the internode communications subsystem.
- Corrects a problem in which a domain panic on the `cluster_root` does not result as it should in a regular panic for the cluster node on which the domain panic occurs.

- Fixes several small issues with `clu_upgrade`:
  - A "process not found" message displayed when finishing the setup stage of `clu_upgrade` has been removed.
  - The ability to roll on a one-node cluster is maintained.
- Addresses a problem on LAN clusters related to improper keep-alive timeouts that can be identified when the following console message is displayed during normal operations (that is, no known failures and no nodes are rebooting):
  - WARNING: ics\_socket\_event: error 60 on channel 0, assume node # is down
- Fixes a problem that occurs when the interconnect is configured using NetRAIN, `cluster_rebuild_delay` is set significantly below the default value, and members are rebooting or failures are occurring on the active links. The console message seen when this occurs is "CNX QDISK: Yielding to foreign owner with provisional quorum."
- Fixes a problem in which I/O barriers may be stalled when a drive becomes hung.
- Prevents write failures from a cluster NFS client that may occur when a second user without write access is concurrently reading the file.
- Fixes a problem that occurs during reboots on heavily loaded cluster using the LAN interconnect and generates the following messages:
  - WARNING: ics\_socket\_event: error 54 on channel 0
  - WARNING: ics\_socket\_event: error 60 on channel 0
- Fixes `kmf` in `drd_kgs_bid_stop_server_io_drained` when a node leaves during a `drd kgs` transaction.
- Corrects a problem in which `drd` continually tries to perform a `munsa unreject` on the drive when a device is deleted while it is in the `munsa reject` state.
- Corrects a problem in which multiple path failures cause `drd` to return `ENODEV` even when a server is available in the cluster.
- Fixes several error handling in `drd` for device error conditions.
- Fixes problem in which a device cannot be opened due to heavy load on the device.
- Fixes a problem in which a CD-ROM is not mountable in a cluster.
- Fixes loss of quorum disk.
- Makes quorum disk parameters configurable.
- Eliminates a window for kernel memory fault panics on AdvFS system calls that are performed via function shipping using the `clu_msfs_syscall_fship` routine.
- Improves `drd` tracing.
- Fixes a Sierra Cluster KCH set free race condition.
- Fixes two errors in `clu_upgrade` that prevents completing the setup stage.
- Prevents a `get_cs_toks()` KMF/assert crash.
- Fixes a `rm_audit_sync_block` panic that occurs when using a long fiber as the Memory Channel interconnect.
- Fixes a timing window in the Internode Communications Subsystem `ddr` device error handling.
- Fixes the `rm_audit_sync_block` panic when using a long fiber with VHUB as the Memory Channel interconnect.
- Fixes `clu_bdmgr` to facilitate CLSM sliced disks for `cluster_root` domain.
- Modifies the manner of checking for user file limits for CFS remote DIO writes.
- Ensures that signals for EFBIG writes are properly generated on a client.
- Ensures the correct processing of CFS in future releases.
- Fixes a multiple free problem of 32-byte memory bucket caused by multiple callbacks from KCH to CLUA.

- Fixes an incorrect if statement, which although a low- risk problem, could block access to a disk device.
- Corrects a confusing error message.
- Fixes a problem seen in a LAN cluster when the CPUs on a member system are not installed contiguously in the lower order slots.
- Allows the quorum disk to be used in spite of transient errors with the quorum disk hardware.
- Corrects an internal logic error that causes the performance of file deletion to be suboptimal.
- Fixes a deadlock that occurs when no members have valid paths to a device and all the nodes in the cluster are attempting failover at the same time.
- Fixes problems seen in the TruCluster RDG component.
- Fixes a race condition in a routine that allocates memory for Memory Channel logical rail and physical rail use. It prevents a KMF during boot, occasionally seen on some AlphaServer GS1280 systems.
- Fixes a race condition which leads to a panic that occurs when a device is deleted on a busy system.
- Adds the ability to log enabled DRD events to circular memory buffer.
- Corrects an Invalid Current Server panic.
- Increases tolerance for intermittent disk boot disk errors early in the boot process.
- Corrects a problem in which I/O operations hang when I/O barriers fail due to the loss of access to drives.
- Fixes a TruCluster NFS server failure that occurs when clients access file systems forcibly removed with the cfsmgr -u command.
- Fixes an incorrect return status for asynchronous direct I/O reads in a cluster if the read request goes beyond the end of the file (EOF).
- Fixes the problem of unintentional loading of gated when nogated is specified with other requested cluamgr operations.
- Fixes a problem in which backplane RAID devices can become inaccessible.
- Provides the following tape-related fixes:
  - Corrects a problem in which hwmgr redirect commands fail on tape devices.
  - Prevents the reuse of a disk number upon deleting and adding a new tape.
  - Corrects a problem in which drdmgr commands can hang on tapes.
  - Updates the code base to make failbacks more proactive.
- Improves defenses against user error during the roll stage of rolling upgrade.
- Fixes TruCluster Distributed Lock Manager (dlm) system panic due to lock transaction ID's being out of synch after a rebuild.
- Corrects a problem in which the TruCluster component DRD (Device Request Dispatcher) does not always return standard error codes.
- Prevents a kernel memory fault panic when drd\_open is called on a device with a valid local path that has no local devt passed in, and this member has the lowest cluster ID of any member in the cluster.
- Prevents CFS token sequence number reuse errors on fast systems.
- Prevent domain panic on a file system that is local to a failed cluster member.
- Prevent CFS write() from updating file access time or panicking on a directory.
- Modifies the way the clu\_upgrade command behaves regarding the availability of backup space in the setup and preinstall stages and adds an appropriate error message.
- Corrects a problem within the TruCluster Kernel Group Services (kgs/kch) subsystem in which the simultaneous booting of multiple nodes may result in a panic due to an unknown node in a remote member node list.
- Removes a delay in the TruCluster component DRD (Device Request Dispatcher) event threads during system booting.

- Corrects a kernel memory fault in `drd_local_device_close`.
- Fixes a kernel memory fault issue on LAN-based clusters that do not have a Memory Channel adapter installed on the systems.
- Fixes problem of non-root users not being able to execute the `caa_stat` command.
- Provides enhancements to CAA commands and the `caad` daemon.
- Resolves a resource exhaustion problem in the TruCluster `kgs/kch` subsystem on high-end clusters, typically with large storage configurations.
- Fixes an assert failure in `cfs` the server.
- Resolves a problem that occurs when adjusting `sysconfig clua` attributes `sticky_entry_timeout` and `sticky_db_cleanup_interval`.
- Ensures that if only a portion of an AIO/DIO write completes, the correct number of bytes written will be returned.
- Allows CFS to correctly handle a token race condition without creating a panic.
- Prevents a node in a cluster from hanging at boot time.
- Corrects misspellings of file system in the `cfsmgr` utility.
- Implements the fast fail policy within DRD.
- Corrects a problem in which backplane RAID devices can become inaccessible when installed on systems running Version 5.1B-2 (Patch Kit 4).
- Enhances the `fuser` command to provide cluster-wide query capability.
- Ensures that the number of `icsmct` receive threads does not exceed the number of CPUs.
- Corrects a condition in which `drd_get_disk_attributes` hang if too many errors are encountered, causing new devices to be inaccessible in a cluster from some cluster members.
- Corrects a problem in which a cluster CFS client would panic in `cfscall_writepages`, reporting `ASSERT (error != EDQUOT)`. This correction eliminates that failure and allows for the proper writing up to the fileset quota and to the end of space for a domain.
- Fixes a rare, three-way deadlock condition when Internode Communication Services (ICS) traffic is in a throttled state and a cluster member that is participating in the throttled traffic is halted.
- Fixes a kernel memory fault in `strlen` on a cluster member during a mount of an AdvFS files system with an improperly specified file system.
- Allows the `ulimit -f` command to function correctly in a cluster.
- Prevents a kernel memory fault panic that may occur with client writes on nearly full domains.
- Prevents a panic on a device close when device connectivity is lost.
- Fixes a mounting KMF of partitioning file system in a cluster.
- Fixes a problem in which a CAA resource and its dependents become inaccessible when the resource fails to start on the node where it is failed over to and there are no more nodes to consider for failover.
- Fixes Oracle socket connection problem.
- Fixes incorrect error handling that could result in memory leak.
- Provides event definitions for traps in cluster MIB files to support Openview NMS.
- Modifies `ics_tcp` to check response buffer for NULL before freeing it.
- Fixes a problem in which booting times in excess of 2 hours occur in a two-node LAN cluster using an `ee (DE6xx)` adapter as the cluster interconnect and connected directly by a crossover cable.
- Corrects a scenario during a cluster member boot whereby a booting member may cause booted members to panic on a kernel memory fault shortly after the messages "Registering CMS Services" and the "rm slave" are printed to the booting console for each MC card.
- Fixes a problem that could cause the system panic "`clua_realloc_port: corrupt list pointers panic`".

- Corrects trapOID for traps generated from the clu\_mibs subagent and provides event definitions for traps in MIB files to support Openview.
- Fixes an inappropriate message that is displayed during CAA resource relocation when invoked from SysMan.
- Fixes 64-byte memory leak in the drd/kgs interface.
- Modifies CNX to check for communication errors while a node joins the cluster.
- Fixes a synchronization issue with a cluster alias ID set among cluster members.
- Prevents a panic from occurring during a failover mount if the AdvFS on-disk file system ID (fsid) does not match the current cluster-wide fsid for the file system.
- Fixes an intermittent core issue in the aliasd daemon caused by improper handling of the interface list.
- Fixes an assertion panic "set-num\_rmt\_mbr\_nodes = 0".
- Prevents a single-node panic in a cluster than can occur under the following conditions:
  - A memory file system of size 4GB or greater is created with the default 512-byte sector size.
  - A memory file system of size 2GB or greater is created with a 1024-byte sector size and other sector sizes.
- Prevents a kernel memory fault panic that may be seen under certain error conditions with MFS file systems.
- Corrects a problem with kch memory usage.

## Patch 27002.00

### *TCRMAN540*

- Provides a new command, clu\_ping, to determine the status of the interconnects in a stretched cluster environment.
- Updates the caa\_relocate(8) and cluamgr(8) reference pages.
- Revises the clua\_services(4) and sys\_attrs\_clua(5) TruCluster reference pages.

---

# 4 Worldwide Language Support Patches

This chapter provides information about the release notes and patches included in Version 5.1B-6 for the Worldwide Language Support (WLS) subset for the Tru64 UNIX operating system.

This chapter is organized as follows:

- “New Release Notes” section lists release notes that are specific to this kit.
- “Summary of Worldwide Language Support Patches” section list the patches that were newly added in this version and the “Patches Delivered in Previous Kits” section describes the patches that were delivered with previous kits.

## New Release Notes

This release contains no defect fixes. There are no new features or enhancements added in this release of Worldwide Language Support software.

## Prior Release Notes

The following sections describe some of the key features and enhancements that were first delivered in previous patch kits.

## Localization of Message Catalog

The messages in the `timezones.msg` for Japanese and Chinese locales is modified to reflect changes in the corresponding English version file that has been updated after the V5.1B-4 release.

The messages in the catalog file is used for a timezone selection.

## Summary of Worldwide Language Support Patches

The following sections provide brief descriptions of the changes delivered in the Version 5.1B-6 patch kit and in previous Version 5.1B patch kits for the WLS subset of the Tru64 UNIX operating system.

Each patch provides fixes to subsets of the WLS software. Subset names (listed in italic font in the following list) consist of four parts; for example, for subset *IOSJPBASE540*, the *IOS* indicates that the subset is part of the WLS package, the *JP* indicates a locale, the *BASE* indicates a category, and the *540* indicates that the subset belongs to the Version 5.1B operating system.

## New Patches

This release contains no defect fixes. There are no new features or enhancements added in this release.

## Patches Delivered in Previous Kits

The following WLS patches were delivered in previous Version 5.1B patch kits. These patches will be installed on your system if you did not install the previous kit.

### Patch 28102.00

*IOSJPBASE540*

- Provides modified localization messages for Japanese and Chinese locales. The modified localization messages are reflections of the English version that have been updated after the V5.1B-4 release.

### Patch 28214.00

*IOSWWBASE540*

- Corrects a problem with `i conv` converter that generates wrong character codes for a Unicode surrogate pair.
- Includes `ICONV_OLD_SURROGATE` environment variable for backward compatibility. When the `ICONV_OLD_SURROGATE` environment variable is set to a non-null value, `i conv` generates the same codes as before.

#### Patch 28256.00

##### *IOSZHCBASE540*

- Provides modified localization messages for Japanese and Chinese locales. The modified localization messages are reflections of the English version that have been updated after the V5.1B-4 release.

#### Patch 27102.00

##### *IOSJPBASE540*

- Provides modified localization messages for Japanese and Chinese locales. The modified localization messages are reflections of the English version ones that have been updated after the V5.1B-3 release.
- Updates some Japanese message catalog files to keep up with corresponding English message catalog files that have been updated after the V5.1B-3 release.
- Updates the Japanese message file to keep up with the corresponding English version that has been updated after the V5.1B-3 release.
- Provides updated Japanese message files for the SysMan suitlets and others.
- Fixes a problem that causes the `i conv` converter to produce incorrect strings every 4,096 bytes of the output in the codeset conversion between Japanese mainframe codesets (`ibmkanji/JEF/KEIS`) and other supported Japanese codesets, such as `euCJP`.
- Fixes an ASU problem in Japanese locales where a double-byte hyphen character in a file name is incorrectly converted to an ASCII underscore character.
- Updates several Japanese message catalog files to catch up with corresponding English message catalog files that have been updated.
- Fixes a problem in which unnecessary quotation marks are included in the Japanese messages for Japanese dictionary's converters, `jsy2vjetxt`, `jsy2wxtxt`, `vje2jsytxt`, and `wx2jsytxt`.
- Updates the Japanese version of `more.cat` to keep it in line with the English version.
- Fixes a problem in which the Japanese dictionary's converters `jsy2vjetxt`, `jsy2wxtxt`, `vje2jsytxt`, and `wx2jsytxt` do not display Japanese messages even though the message catalog files are provided.
- Updates some Japanese message catalog files to keep up with corresponding English message catalog files that have been updated after the V5.1B-2 release.
- Updates some Japanese message catalog files to keep up with corresponding English message catalog files that have been updated after the V5.1B-2 release.

#### Patch 27103.00

##### *IOSJPBIN540*

- Fixes a problem with the Asian `tty` that can cause a kernel crash and a hang-up problem during a certain type of the login/logout stress test.

#### Patch 27106.00

##### *IOSJPCDEDT540*

- Updates some Japanese message catalog files to keep up with corresponding English message catalog files.
- Corrects a potential security vulnerability where, under certain circumstances, system integrity may be compromised. This may be in the form of improper file access.



- Provides new error messages for the dtprintinfo and dtwm applications.
- Makes applications display correct error messages under the Japanese locale setting.
- Updates the Japanese version of message and uil files to keep up with the English version.

#### Patch 27125.00

##### *IOSJPSYSMAN540*

- Adds localization messages for Japanese locales. The messages are to be added to catch up with the addition of messages in the English version that have been updated after the V5.1B-3 release.
- Provides new Japanese online help files for SysMan file system applications as well as update translations for other SysMan applications.
- Provides updated Japanese message files for the SysMan suitlets and others.
- Updates the Japanese version of message and uil files to keep up with the English version.
- Updates some Japanese message catalog files to keep up with corresponding English message catalog files which have been updated after the V5.1B-2 release.

#### Patch 27131.00

##### *IOSJPXADMIN540*

- Provides updated Japanese message files for the SysMan suitlets and others.

#### Patch 27134.00

##### *IOSJPXSYSMAN540*

- Provides new Japanese online help files for SysMan file system applications as well as update translations for other SysMan applications.
- Provides updated Japanese message files for the SysMan suitlets and others.
- Updates the Japanese version of message and uil files to keep up with the English version.

#### Patch 27214.00

*IOSWWBASE540*

- Fixes a display width mismatch problem in the zh\_CN.GB18030 locale.

#### Patch 27215.00

*IOSWWBIN540*

- Fixes a problem with the Asian tty that can cause a kernel crash and a hang-up problem during a certain type of the login/logout stress test.

#### Patch 27248.00

*IOSWWSYSMAN540*

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

#### Patch 27251.00

*IOSWWX11540*

- Fixes a problem in which the Qu-Wei input method invoked from dxhanziim or dxim does not produce a correct character from Unicode value.

#### Patch 27256.00

*IOSZHCNBASE540*

- Provides modified localization messages for Japanese and Chinese locales. The modified localization messages are reflections of the English version ones that have been updated after the V5.1B-3 release.
- Fixes a display width mismatch problem in the zh\_CN.GB18030 locale.

#### Patch 27271.00

*IOSZHSSYSMAN540*

- Provides an updated Chinese message file for the SysMan applications and others.

#### Patch 27276.00

*IOSZHSXADMIN540*

- Provides an updated Chinese message file for the SysMan applications and others.

---

## 5 CSPs Included in This Kit

This chapter lists the customer-specific patches (CSPs) for Tru64 UNIX and TruCluster Server that are superseded by patches in this kit.

To find out which CSPs are on your system, use the `dupatch` utility's patch tracking feature, as described in the [Patch Kit Installation Instructions](#). If you are installing patches from the command line, see the appendix "Using `dupatch` from the Command Line" in the [Patch Kit Installation Instructions](#) for information about determining which CSPs are on your system.

When using the Patch Tracking feature, `dupatch` displays either "Tru64\_UNIX\_V5.1B" or "TruCluster\_V5.1B" before the patch number, and also displays the type of patch, such as Security related or Network; for example:

- Tru64\_UNIX\_V5.1B / Commands, Shells, & Utilities Patches:  
Patch C 00021.00 - ksh processes use excessive CPU time
- TruCluster\_V5.1B / Cluster Filesystem Patches:  
Patch C 00005.00 - CFS proplist access corrections

To see if the functionality delivered in a CSP is included in this kit or previous V5.1B patch kits, note the number returned in the previous example and check to see if that number is listed in the following tables.

Note that several cluster file system CSPs are included in the Tru64 UNIX list because they are dependent on file system patches for the operating system.

## New Superseded Tru64 UNIX CSPs

The numbers in the following tables represent Tru64 UNIX CSPs that are superseded in the current release.

**Table 5-1 Superseded CSPs 02298.01 to 02493.00**

|                  |                  |                  |                  |
|------------------|------------------|------------------|------------------|
| Patch C 02298.01 | Patch C 02371.01 | Patch C 02426.01 | Patch C 02447.00 |
| Patch C 02356.00 | Patch C 02402.00 | Patch C 02432.00 | Patch C 02450.00 |
| Patch C 02357.00 | Patch C 02410.01 | Patch C 02433.00 | Patch C 02452.01 |
| Patch C 02360.00 | Patch C 02419.00 | Patch C 02438.01 | Patch C 02467.00 |
| Patch C 02363.00 | Patch C 02420.00 | Patch C 02440.00 | Patch C 02470.00 |
| Patch C 02366.00 | Patch C 02422.00 | Patch C 02441.00 | Patch C 02471.00 |
| Patch C 02370.01 | Patch C 02426.00 | Patch C 02444.00 | Patch C 02492.00 |
|                  |                  |                  | Patch C 02493.00 |

## Tru64 UNIX CSPs Superseded in Previous Kits

The numbers in the following tables represent Tru64 UNIX CSPs that were included in previous kits.

**Table 5-2 Superseded CSPs 01936.01 to 02355.00**

|                  |                  |                  |                  |
|------------------|------------------|------------------|------------------|
| Patch C 01936.01 | Patch C 02080.00 | Patch C 02202.00 | Patch C 02271.00 |
| Patch C 01937.00 | Patch C 02081.00 | Patch C 02205.00 | Patch C 02273.00 |
| Patch C 01944.04 | Patch C 02083.00 | Patch C 02209.01 | Patch C 02281.00 |
| Patch C 01945.02 | Patch C 02084.00 | Patch C 02213.02 | Patch C 02285.00 |
| Patch C 01947.00 | Patch C 02091.00 | Patch C 02214.00 | Patch C 02286.00 |
| Patch C 01949.00 | Patch C 02092.00 | Patch C02222.01  | Patch C 02287.01 |
| Patch C 01956.00 | Patch C 02093.00 | Patch C 02225.00 | Patch C 02288.00 |
| Patch C 01960.00 | Patch C 02094.00 | Patch C 02230.00 | Patch C 02301.00 |
| Patch C 01977.00 | Patch C 02096.00 | Patch C 02234.00 | Patch C 02305.00 |
| Patch C 01988.00 | Patch C 02097.00 | Patch C 02235.00 | Patch C 02305.01 |
| Patch C 01990.00 | Patch C 02105.00 | Patch C 02236.00 | Patch C 02307.00 |
| Patch C 02002.00 | Patch C 02109.00 | Patch C 02237.00 | Patch C 02309.00 |
| Patch C 02016.01 | Patch C 02127.00 | Patch C 02238.00 | Patch C 02310.00 |
| Patch C 02017.00 | Patch C 02129.00 | Patch C 02239.00 | Patch C 02312.00 |
| Patch C 02023.00 | Patch C 02130.00 | Patch C 02240.00 | Patch C 02313.00 |
| Patch C 02023.01 | Patch C 02140.00 | Patch C 02241.00 | Patch C 02314.00 |
| Patch C 02026.00 | Patch C 02144.00 | Patch C 02245.02 | Patch C 02315.00 |
| Patch C 02028.00 | Patch C 02145.00 | Patch C 02247.00 | Patch C 02321.00 |
| Patch C 02029.00 | Patch C 02147.00 | Patch C 02248.00 | Patch C 02324.00 |
| Patch C 02030.00 | Patch C 02149.00 | Patch C 02249.00 | Patch C 02325.00 |
| Patch C 02039.01 | Patch C 02151.00 | Patch C 02251.00 | Patch C 02326.00 |
| Patch C 02041.00 | Patch C 02152.00 | Patch C02251.01  | Patch C 02326.01 |
| Patch C 02048.00 | Patch C 02153.00 | Patch C 02252.00 | Patch C 02328.00 |
| Patch C 02054.00 | Patch C 02154.00 | Patch C 02253.00 | Patch C 02330.00 |
| Patch C 02060.00 | Patch C 02155.00 | Patch C 02254.00 | Patch C 02330.01 |
| Patch C 02062.00 | Patch C 02156.00 | Patch C 02255.00 | Patch C 02330.02 |
| Patch C 02063.00 | Patch C 02157.02 | Patch C 02256.02 | Patch C 02339.00 |
| Patch C 02065.00 | Patch C 02163.00 | Patch C 02258.01 | Patch C 02340.00 |
| Patch C 02067.01 | Patch C 02164.00 | Patch C 02260.00 | Patch C 02341.00 |
| Patch C 02068.01 | Patch C 02164.01 | Patch C 02260.01 | Patch C 02342.00 |
| Patch C 02069.00 | Patch C 02165.00 | Patch C 02261.00 | Patch C 02343.00 |
| Patch C 02070.00 | Patch C 02171.02 | Patch C 02263.00 | Patch C 02344.00 |
| Patch C 02071.00 | Patch C 02172.00 | Patch C 02264.00 | Patch C 02349.01 |
| Patch C 02072.00 | Patch C 02176.00 | Patch C 02265.00 | Patch C 02351.00 |
| Patch C 02073.00 | Patch C 02179.00 | Patch C 02267.00 | Patch C 02352.00 |
| Patch C 02073.01 | Patch C 02179.01 | Patch C 02268.00 | Patch C 02354.00 |
| Patch C 02079.00 | Patch C 02201.01 | Patch C 02270.00 | Patch C 02355.00 |

**Table 5-3 Superseded CSPs 00021.00 to 00283.00**

|                  |                  |                  |                  |
|------------------|------------------|------------------|------------------|
| Patch C 00021.00 | Patch C 00110.00 | Patch C 00159.01 | Patch C 00220.00 |
| Patch C 00027.00 | Patch C 00112.00 | Patch C 00160.00 | Patch C 00221.00 |
| Patch C 00028.00 | Patch C 00112.01 | Patch C 00161.00 | Patch C 00222.00 |
| Patch C 00029.00 | Patch C 00113.00 | Patch C 00162.00 | Patch C 00223.00 |
| Patch C 00030.00 | Patch C 00115.00 | Patch C 00163.00 | Patch C 00224.00 |
| Patch C 00031.02 | Patch C 00115.01 | Patch C 00164.00 | Patch C 00225.00 |
| Patch C 00034.00 | Patch C 00115.02 | Patch C 00165.00 | Patch C 00226.00 |
| Patch C 00037.00 | Patch C 00118.00 | Patch C 00166.00 | Patch C 00227.00 |
| Patch C 00039.00 | Patch C 00122.00 | Patch C 00166.01 | Patch C 00229.00 |
| Patch C 00041.00 | Patch C 00123.00 | Patch C 00166.02 | Patch C 00230.01 |
| Patch C 00042.00 | Patch C 00123.01 | Patch C 00168.00 | Patch C 00231.02 |
| Patch C 00044.00 | Patch C 00124.00 | Patch C 00170.00 | Patch C 00232.00 |
| Patch C 00051.00 | Patch C 00124.01 | Patch C 00171.00 | Patch C 00233.00 |
| Patch C 00052.00 | Patch C 00125.00 | Patch C 00172.00 | Patch C 00234.00 |
| Patch C 00062.00 | Patch C 00136.00 | Patch C 00174.00 | Patch C 00235.00 |
| Patch C 00063.00 | Patch C 00136.01 | Patch C 00175.00 | Patch C 00237.00 |
| Patch C 00066.00 | Patch C 00137.00 | Patch C 00176.00 | Patch C 00239.00 |
| Patch C 00067.00 | Patch C 00137.01 | Patch C 00176.01 | Patch C 00251.00 |
| Patch C 00073.00 | Patch C 00137.02 | Patch C 00177.00 | Patch C 00253.00 |
| Patch C 00074.00 | Patch C 00138.00 | Patch C 00177.01 | Patch C 00253.01 |
| Patch C 00075.00 | Patch C 00138.01 | Patch C 00179.00 | Patch C 00255.00 |
| Patch C 00075.01 | Patch C 00139.00 | Patch C 00180.00 | Patch C 00256.00 |
| Patch C 00075.02 | Patch C 00144.00 | Patch C 00181.00 | Patch C 00257.00 |
| Patch C 00075.03 | Patch C 00144.01 | Patch C 00183.00 | Patch C 00262.00 |
| Patch C 00076.00 | Patch C 00145.00 | Patch C 00188.00 | Patch C 00263.00 |
| Patch C 00077.00 | Patch C 00146.00 | Patch C 00188.01 | Patch C 00264.00 |
| Patch C 00078.00 | Patch C 00146.01 | Patch C 00189.00 | Patch C 00266.00 |
| Patch C 00080.00 | Patch C 00146.02 | Patch C 00191.00 | Patch C 0266.01  |
| Patch C 00081.00 | Patch C 00147.00 | Patch C 00192.00 | Patch C 00267.00 |
| Patch C 00082.00 | Patch C 00147.01 | Patch C 00195.00 | Patch C 00269.00 |
| Patch C 00086.00 | Patch C 00147.02 | Patch C 00195.01 | Patch C 00270.00 |
| Patch C 00089.00 | Patch C 00148.00 | Patch C 00197.00 | Patch C 00270.01 |
| Patch C 00092.00 | Patch C 00148.01 | Patch C 00199.00 | Patch C 00271.00 |
| Patch C 00092.01 | Patch C 00148.02 | Patch C 00200.00 | Patch C 00272.00 |
| Patch C 00094.00 | Patch C 00149.00 | Patch C 00201.00 | Patch C 00273.00 |
| Patch C 0094.01  | Patch C 00149.01 | Patch C 00201.01 | Patch C 00275.00 |
| Patch C 00095.00 | Patch C 00149.02 | Patch C 00203.00 | Patch C 00276.00 |
| Patch C 00096.00 | Patch C 00150.00 | Patch C 00205.00 | Patch C 00277.01 |
| Patch C 00097.00 | Patch C 00150.01 | Patch C 00208.00 | Patch C 00278.00 |
| Patch C 00098.00 | Patch C 00150.02 | Patch C 00212.00 | Patch C 00279.00 |
| Patch C 00099.00 | Patch C 00151.00 | Patch C 00213.00 | Patch C 00280.00 |
| Patch C 00101.00 | Patch C 00152.00 | Patch C 00214.00 | Patch C 00280.01 |
| Patch C 00104.00 | Patch C 00153.00 | Patch C 00217.00 | Patch C 00281.00 |
| Patch C 00104.01 | Patch C 00158.01 | Patch C 00218.00 | Patch C 00283.00 |
| Patch C 00106.00 | Patch C 00159.00 | Patch C 00219.00 |                  |

**Table 5-4 Superseded CSPs 00283.01 to 00543.00**

|                  |                  |                  |                  |
|------------------|------------------|------------------|------------------|
| Patch C 00283.01 | Patch C 00331.01 | Patch C 00409.01 | Patch C 00476.00 |
| Patch C 00287.00 | Patch C 00332.00 | Patch C 00411.00 | Patch C 00476.01 |
| Patch C 00287.01 | Patch C 00332.01 | Patch C 00414.00 | Patch C 00477.00 |
| Patch C 00288.00 | Patch C 00332.02 | Patch C 00414.01 | Patch C 00477.01 |
| Patch C 00290.00 | Patch C 00333.00 | Patch C 00414.03 | Patch C 00477.02 |
| Patch C 00290.01 | Patch C 00333.01 | Patch C 00415.00 | Patch C 00477.03 |
| Patch C 00291.00 | Patch C 00334.00 | Patch C 00416.00 | Patch C 00477.04 |
| Patch C 00292.00 | Patch C 00337.00 | Patch C 00417.00 | Patch C 00478.00 |
| Patch C 00295.00 | Patch C 00337.02 | Patch C 00419.00 | Patch C 00479.00 |
| Patch C 00296.00 | Patch C 00339.00 | Patch C 00421.01 | Patch C 00480.00 |
| Patch C 00297.00 | Patch C 00343.00 | Patch C 00422.00 | Patch C 00482.00 |
| Patch C 00298.00 | Patch C 00344.00 | Patch C 00423.00 | Patch C 00483.00 |
| Patch C 00299.00 | Patch C 00345.00 | Patch C 00424.00 | Patch C 00484.00 |
| Patch C 00300.00 | Patch C 00346.00 | Patch C 00424.01 | Patch C 00486.00 |
| Patch C 00300.02 | Patch C 00348.00 | Patch C 00424.02 | Patch C 00489.00 |
| Patch C 00300.03 | Patch C 00351.00 | Patch C 00424.03 | Patch C 00490.00 |
| Patch C 00300.04 | Patch C 00351.01 | Patch C 00425.00 | Patch C 00490.01 |
| Patch C 00302.00 | Patch C 00355.00 | Patch C 00426.00 | Patch C 00491.00 |
| Patch C 00302.01 | Patch C 00357.00 | Patch C 00428.00 | Patch C 00496.00 |
| Patch C 00302.02 | Patch C 00357.01 | Patch C 00429.00 | Patch C 00497.00 |
| Patch C 00302.03 | Patch C 00362.00 | Patch C 00430.00 | Patch C 00499.00 |
| Patch C 00303.00 | Patch C 00362.01 | Patch C 00431.00 | Patch C 00500.00 |
| Patch C 00307.00 | Patch C 00362.02 | Patch C 00431.01 | Patch C 00505.00 |
| Patch C 00308.00 | Patch C 00363.00 | Patch C 00432.00 | Patch C 00509.00 |
| Patch C 00308.01 | Patch C 00364.00 | Patch C 00433.00 | Patch C 00510.00 |
| Patch C 00309.03 | Patch C 00369.00 | Patch C 00434.00 | Patch C 00514.00 |
| Patch C 00311.00 | Patch C 00370.00 | Patch C 00434.01 | Patch C 00514.01 |
| Patch C 00312.00 | Patch C 00371.00 | Patch C 00434.02 | Patch C 00516.00 |
| Patch C 00312.01 | Patch C 00381.00 | Patch C 00435.00 | Patch C 00521.00 |
| Patch C 00312.02 | Patch C 00381.01 | Patch C 00441.00 | Patch C 00521.01 |
| Patch C 00312.04 | Patch C 00381.02 | Patch C 00443.00 | Patch C 00522.00 |
| Patch C 00315.00 | Patch C 00382.00 | Patch C 00445.00 | Patch C 00528.00 |
| Patch C 00316.00 | Patch C 00383.00 | Patch C 00446.00 | Patch C 00528.01 |
| Patch C 00319.00 | Patch C 00384.00 | Patch C 00448.00 | Patch C 00529.00 |
| Patch C 00321.00 | Patch C 00386.00 | Patch C 00453.00 | Patch C 00530.00 |
| Patch C 00322.00 | Patch C 00390.00 | Patch C 00458.00 | Patch C 00531.00 |
| Patch C 00323.00 | Patch C 00391.00 | Patch C 00461.00 | Patch C 00113.00 |
| Patch C 00325.00 | Patch C 00393.00 | Patch C 00462.00 | Patch C 00533.00 |
| Patch C 00326.00 | Patch C 00395.00 | Patch C 00464.00 | Patch C 00533.02 |
| Patch C 00327.00 | Patch C 00396.00 | Patch C 00465.01 | Patch C 00538.00 |
| Patch C 00328.00 | Patch C 00399.00 | Patch C 00467.00 | Patch C 00539.00 |
| Patch C 00329.00 | Patch C 00399.01 | Patch C 00469.00 | Patch C 00540.00 |
| Patch C 00330.00 | Patch C 00399.02 | Patch C 00470.00 | Patch C 00541.00 |
| Patch C 00330.03 | Patch C 00407.00 | Patch C 00470.01 | Patch C 00542.00 |
| Patch C 00331.00 | Patch C 00407.01 | Patch C 00471.00 | Patch C 00543.00 |

**Table 5-5 Superseded CSPs 00545.00 to 00855.01**

|                  |                  |                  |                  |
|------------------|------------------|------------------|------------------|
| Patch C 00545.00 | Patch C 00659.00 | Patch C 00752.00 | Patch C 00805.00 |
| Patch C 00548.00 | Patch C 00659.01 | Patch C 00755.01 | Patch C 00805.01 |
| Patch C 00549.00 | Patch C 00668.00 | Patch C 00756.00 | Patch C 00805.04 |
| Patch C 00552.00 | Patch C 00670.05 | Patch C 00757.00 | Patch C 00805.05 |
| Patch C 00553.00 | Patch C 00672.00 | Patch C 00757.01 | Patch C 00806.00 |
| Patch C 00555.00 | Patch C 00673.00 | Patch C 00757.02 | Patch C 00807.00 |
| Patch C 00556.00 | Patch C 00673.01 | Patch C 00757.03 | Patch C 00808.00 |
| Patch C 00556.01 | Patch C 00675.00 | Patch C 00757.04 | Patch C 00809.00 |
| Patch C 00561.00 | Patch C 00677.02 | Patch C 00758.00 | Patch C 00811.00 |
| Patch C 00562.00 | Patch C 00678.00 | Patch C 00761.00 | Patch C 00812.00 |
| Patch C 00563.00 | Patch C 00686.00 | Patch C 00762.00 | Patch C 00815.00 |
| Patch C 00563.01 | Patch C 00687.00 | Patch C 00765.00 | Patch C 00818.00 |
| Patch C 00564.00 | Patch C 00688.00 | Patch C 00766.00 | Patch C 00819.00 |
| Patch C 00564.01 | Patch C 00697.00 | Patch C 00767.00 | Patch C 00822.00 |
| Patch C 00564.02 | Patch C 00699.00 | Patch C 00770.00 | Patch C 00823.00 |
| Patch C 00567.00 | Patch C 00708.00 | Patch C 00771.00 | Patch C 00824.00 |
| Patch C 00570.00 | Patch C 00709.00 | Patch C 00771.01 | Patch C 00825.00 |
| Patch C 00575.00 | Patch C 00711.00 | Patch C 00772.00 | Patch C 00825.01 |
| Patch C 00583.00 | Patch C 00711.01 | Patch C 00773.00 | Patch C 00829.00 |
| Patch C 00586.00 | Patch C 00722.00 | Patch C 00774.00 | Patch C 00830.00 |
| Patch C 00594.00 | Patch C 00722.01 | Patch C 00774.01 | Patch C 00830.01 |
| Patch C 00599.00 | Patch C 00724.00 | Patch C 00776.00 | Patch C 00830.03 |
| Patch C 00600.00 | Patch C 00724.01 | Patch C 00777.00 | Patch C 00832.00 |
| Patch C 00602.00 | Patch C 00724.02 | Patch C 00777.01 | Patch C 00835.00 |
| Patch C 00603.00 | Patch C 00729.00 | Patch C 00782.00 | Patch C 00837.00 |
| Patch C 00607.00 | Patch C 00733.00 | Patch C 00783.00 | Patch C 00837.01 |
| Patch C 00608.00 | Patch C 00736.00 | Patch C 00783.01 | Patch C 00837.02 |
| Patch C 00609.00 | Patch C 00736.01 | Patch C 00783.02 | Patch C 00838.00 |
| Patch C 00609.01 | Patch C 00736.02 | Patch C 00784.00 | Patch C 00839.00 |
| Patch C 00612.00 | Patch C 00738.00 | Patch C 00785.00 | Patch C 00840.00 |
| Patch C 00613.00 | Patch C 00738.01 | Patch C 00785.01 | Patch C 00840.04 |
| Patch C 00613.01 | Patch C 00740.00 | Patch C 00785.02 | Patch C 00840.05 |
| Patch C 00618.00 | Patch C 00741.00 | Patch C 00785.03 | Patch C 00843.00 |
| Patch C 00619.00 | Patch C 00742.00 | Patch C 00786.00 | Patch C 00846.00 |
| Patch C 00620.00 | Patch C 00742.01 | Patch C 00791.00 | Patch C 00846.01 |
| Patch C 00622.00 | Patch C 00743.00 | Patch C 00795.00 | Patch C 00849.00 |
| Patch C 00626.00 | Patch C 00743.01 | Patch C 00795.02 | Patch C 00849.01 |
| Patch C 00627.00 | Patch C 00743.02 | Patch C 00797.00 | Patch C 00849.05 |
| Patch C 00628.00 | Patch C 00744.04 | Patch C 00798.00 | Patch C 00849.07 |
| Patch C 00630.00 | Patch C 00745.04 | Patch C 00799.00 | Patch C 00849.08 |
| Patch C 00636.00 | Patch C 00750.00 | Patch C 00800.00 | Patch C 00849.09 |
| Patch C 00639.00 | Patch C 00750.01 | Patch C 00801.00 | Patch C 00849.10 |
| Patch C 00639.01 | Patch C 00750.02 | Patch C 00802.00 | Patch C 00849.11 |
| Patch C 00642.00 | Patch C 00750.03 | Patch C 00803.00 | Patch C 00850.00 |
| Patch C 00644.00 | Patch C 00750.04 | Patch C 00803.01 | Patch C 00855.01 |
| Patch C 00658.00 | Patch C 00751.00 | Patch C 00804.00 |                  |



**Table 5-6 Superseded CSPs 00856.00 to 01107.00**

|                  |                  |                  |                  |
|------------------|------------------|------------------|------------------|
| Patch C 00856.00 | Patch C 00916.01 | Patch C 00980.00 | Patch C 01042.01 |
| Patch C 00856.01 | Patch C 00917.00 | Patch C 00982.00 | Patch C 01047.00 |
| Patch C 00858.00 | Patch C 00920.00 | Patch C 00983.00 | Patch C 01048.01 |
| Patch C 00859.00 | Patch C 00923.00 | Patch C 00983.01 | Patch C 01048.02 |
| Patch C 00861.00 | Patch C 00924.00 | Patch C 00985.00 | Patch C 01050.00 |
| Patch C 00863.00 | Patch C 00925.00 | Patch C 00986.00 | Patch C 01051.00 |
| Patch C 00865.00 | Patch C 00926.00 | Patch C 00987.00 | Patch C 01052.00 |
| Patch C 00867.00 | Patch C 00927.00 | Patch C 00989.00 | Patch C 01055.00 |
| Patch C 00867.02 | Patch C 00929.00 | Patch C 00991.00 | Patch C 01056.00 |
| Patch C 00868.00 | Patch C 00932.00 | Patch C 00992.00 | Patch C 01057.00 |
| Patch C 00869.00 | Patch C 00932.01 | Patch C 00993.0  | Patch C 01058.00 |
| Patch C 00872.05 | Patch C 00933.00 | Patch C 00994.00 | Patch C 01062.00 |
| Patch C 00875.03 | Patch C 00933.03 | Patch C 00994.02 | Patch C 01063.00 |
| Patch C 00876.0  | Patch C 00933.04 | Patch C 00996.00 | Patch C 01067.00 |
| Patch C 00878.03 | Patch C 00933.01 | Patch C 00997.00 | Patch C 01068.00 |
| Patch C 00880.00 | Patch C 00935.00 | Patch C 01000.00 | Patch C 01069.00 |
| Patch C 00881.00 | Patch C 00938.02 | Patch C 01002.00 | Patch C 01070.00 |
| Patch C 00882.00 | Patch C 00939.00 | Patch C 01003.00 | Patch C 01070.01 |
| Patch C 00883.00 | Patch C 00939.01 | Patch C 01004.00 | Patch C 01070.02 |
| Patch C 00884.00 | Patch C 00941.00 | Patch C 01004.01 | Patch C 01070.03 |
| Patch C 00884.01 | Patch C 00941.01 | Patch C 01006.00 | Patch C 01070.04 |
| Patch C 00885.00 | Patch C 00942.00 | Patch C 01007.00 | Patch C 01070.05 |
| Patch C 00885.01 | Patch C 00952.00 | Patch C 01007.01 | Patch C 01070.06 |
| Patch C 00886.00 | Patch C 00953.00 | Patch C 01008.00 | Patch C 01070.07 |
| Patch C 00887.00 | Patch C 00954.00 | Patch C 01010.00 | Patch C 01071.00 |
| Patch C 00888.00 | Patch C 00956.00 | Patch C 01011.00 | Patch C 01072.00 |
| Patch C 00889.00 | Patch C 00958.00 | Patch C 01011.01 | Patch C 01073.00 |
| Patch C 00892.00 | Patch C 00958.01 | Patch C 01013.00 | Patch C 01079.00 |
| Patch C 00892.01 | Patch C 00959.00 | Patch C 01013.01 | Patch C 01079.01 |
| Patch C 00896.00 | Patch C 00960.00 | Patch C 01014.00 | Patch C 01081.00 |
| Patch C 00898.00 | Patch C 00960.01 | Patch C 01016.00 | Patch C 01081.01 |
| Patch C 00899.00 | Patch C 00961.00 | Patch C 01025.00 | Patch C 01082.00 |
| Patch C 00900.00 | Patch C 00961.01 | Patch C 01026.00 | Patch C 01085.00 |
| Patch C 00900.01 | Patch C 00963.00 | Patch C 01027.00 | Patch C 01088.00 |
| Patch C 00903.00 | Patch C 00964.00 | Patch C 01030.00 | Patch C 01095.00 |
| Patch C 00903.01 | Patch C 00964.02 | Patch C 01032.00 | Patch C 01098.00 |
| Patch C 00905.00 | Patch C 00965.00 | Patch C 01036.00 | Patch C 01100.00 |
| Patch C 00905.01 | Patch C 00966.00 | Patch C 01036.01 | Patch C 01102.00 |
| Patch C 00906.00 | Patch C 00971.00 | Patch C 01036.03 | Patch C 01103.00 |
| Patch C 00907.00 | Patch C 00973.00 | Patch C 01036.04 | Patch C 01105.00 |
| Patch C 00908.00 | Patch C 00974.00 | Patch C 01036.05 | Patch C 01107.00 |
| Patch C 00909.00 | Patch C 00975.00 | Patch C 01036.06 |                  |
| Patch C 00914.01 | Patch C 00979.00 | Patch C 01041.00 |                  |
| Patch C 00916.00 | Patch C 00975.00 | Patch C 01042.00 |                  |

**Table 5-7 Superseded CSPs 00538.01 to 01361.02**

|                  |                  |                  |                  |
|------------------|------------------|------------------|------------------|
| Patch C 00538.01 | Patch C 01151.01 | Patch C 01213.00 | Patch C 01283.00 |
| Patch C 00833.00 | Patch C 01155.00 | Patch C 01214.00 | Patch C 01284.00 |
| Patch C 00933.02 | Patch C 01155.01 | Patch C 01214.01 | Patch C 01288.00 |
| Patch C 00999.00 | Patch C 01156.00 | Patch C 01214.02 | Patch C 01289.00 |
| Patch C 01018.00 | Patch C 01159.00 | Patch C 01217.00 | Patch C 01292.00 |
| Patch C 01020.00 | Patch C 01163.00 | Patch C 01219.00 | Patch C 01293.00 |
| Patch C 01028.00 | Patch C 01164.00 | Patch C 01220.01 | Patch C 01294.00 |
| Patch C 01035.00 | Patch C 01173.00 | Patch C 01222.00 | Patch C 01299.00 |
| Patch C 01039.00 | Patch C 01174.00 | Patch C 01223.00 | Patch C 01300.00 |
| Patch C 01046.00 | Patch C 01175.00 | Patch C 01226.00 | Patch C 01302.00 |
| Patch C 01053.00 | Patch C 01179.00 | Patch C 01230.00 | Patch C 01303.00 |
| Patch C 01080.00 | Patch C 01181.02 | Patch C 01231.00 | Patch C 01306.00 |
| Patch C 01087.00 | Patch C 01181.03 | Patch C 01232.00 | Patch C 01307.00 |
| Patch C 01094.00 | Patch C 01182.00 | Patch C 01233.01 | Patch C 01314.00 |
| Patch C 01094.01 | Patch C 01184.00 | Patch C 01235.00 | Patch C 01316.00 |
| Patch C 01101.00 | Patch C 01185.00 | Patch C 01237.01 | Patch C 01319.00 |
| Patch C 01101.01 | Patch C 01186.00 | Patch C 01238.00 | Patch C 01323.00 |
| Patch C 01114.01 | Patch C 01187.00 | Patch C 01239.00 | Patch C 01323.01 |
| Patch C 01114.02 | Patch C 01188.00 | Patch C 01241.03 | Patch C 01327.00 |
| Patch C 01114.03 | Patch C 01190.00 | Patch C 01244.00 | Patch C 01329.00 |
| Patch C 01121.01 | Patch C 01196.00 | Patch C 01244.01 | Patch C 01329.01 |
| Patch C 01130.01 | Patch C 01197.00 | Patch C 01249.00 | Patch C 01332.00 |
| Patch C 01132.00 | Patch C 01198.00 | Patch C 01250.00 | Patch C 01333.00 |
| Patch C 01133.00 | Patch C 01199.00 | Patch C 01252.00 | Patch C 01334.03 |
| Patch C 01134.00 | Patch C 01199.01 | Patch C 01253.01 | Patch C 01335.00 |
| Patch C 01137.00 | Patch C 01199.02 | Patch C 01253.02 | Patch C 01345.00 |
| Patch C 01138.00 | Patch C 01199.03 | Patch C 01254.00 | Patch C 01348.00 |
| Patch C 01139.00 | Patch C 01199.04 | Patch C 01259.00 | Patch C 01348.01 |
| Patch C 01141.00 | Patch C 01200.00 | Patch C 01261.00 | Patch C 01352.00 |
| Patch C 01142.00 | Patch C 01200.01 | Patch C 01261.01 | Patch C 01353.00 |
| Patch C 01142.01 | Patch C 01202.00 | Patch C 01262.00 | Patch C 01354.00 |
| Patch C 01142.02 | Patch C 01202.01 | Patch C 01265.00 | Patch C 01354.01 |
| Patch C 01142.04 | Patch C 01204.00 | Patch C 01268.00 | Patch C 01354.02 |
| Patch C 01143.00 | Patch C 01205.00 | Patch C 01271.00 | Patch C 01354.03 |
| Patch C 01145.00 | Patch C 01205.02 | Patch C 01273.00 | Patch C 01354.04 |
| Patch C 01146.00 | Patch C 01207.00 | Patch C 01276.00 | Patch C 01358.08 |
| Patch C 01146.01 | Patch C 01209.00 | Patch C 01277.00 | Patch C 01360.00 |
| Patch C 01147.00 | Patch C 01210.00 | Patch C 01279.00 | Patch C 01361.00 |
| Patch C 01147.01 | Patch C 01211.00 | Patch C 01281.00 | Patch C 01361.01 |
| Patch C 01148.00 | Patch C 01212.01 | Patch C 01282.00 | Patch C 01361.02 |

**Table 5-8 Superseded CSPs 01362.00 to Patch C 01934.00**

|                  |                  |                  |                  |
|------------------|------------------|------------------|------------------|
| Patch C 01362.00 | Patch C 01456.04 | Patch C 01584.00 | Patch C 01781.00 |
| Patch C 01362.02 | Patch C 01457.00 | Patch C 01587.00 | Patch C 01781.01 |
| Patch C 01363.04 | Patch C 01457.01 | Patch C 01590.00 | Patch C 01783.00 |
| Patch C 01375.00 | Patch C 01457.02 | Patch C 01592.00 | Patch C 01786.00 |
| Patch C 01376.00 | Patch C 01463.00 | Patch C 01594.00 | Patch C 01787.00 |
| Patch C 01379.00 | Patch C 01464.00 | Patch C 01600.00 | Patch C 01790.00 |
| Patch C 01380.00 | Patch C 01465.00 | Patch C 01608.00 | Patch C 01790.01 |
| Patch C 01380.01 | Patch C 01466.00 | Patch C 01610.00 | Patch C 01790.02 |
| Patch C 01382.00 | Patch C 01467.00 | Patch C 01611.00 | Patch C 01790.03 |
| Patch C 01383.00 | Patch C 01468.00 | Patch C 01614.00 | Patch C 01790.04 |
| Patch C 01384.00 | Patch C 01468.01 | Patch C 01614.01 | Patch C 01803.00 |
| Patch C 01385.00 | Patch C 01468.02 | Patch C 01618.00 | Patch C 01804.00 |
| Patch C 01387.00 | Patch C 01471.00 | Patch C 01620.00 | Patch C 01806.00 |
| Patch C 01388.00 | Patch C 01471.01 | Patch C 01621.00 | Patch C 01813.00 |
| Patch C 01388.02 | Patch C 01472.00 | Patch C 01622.00 | Patch C 01816.00 |
| Patch C 01389.00 | Patch C 01472.01 | Patch C 01624.00 | Patch C 01818.00 |
| Patch C 01389.06 | Patch C 01472.02 | Patch C 01626.00 | Patch C 01818.01 |
| Patch C 01389.07 | Patch C 01472.03 | Patch C 01627.00 | Patch C 01826.00 |
| Patch C 01389.10 | Patch C 01473.00 | Patch C 01630.00 | Patch C 01827.00 |
| Patch C 01389.11 | Patch C 01474.00 | Patch C 01631.00 | Patch C 01828.00 |
| Patch C 01389.12 | Patch C 01474.01 | Patch C 01636.00 | Patch C 01829.00 |
| Patch C 01389.15 | Patch C 01474.02 | Patch C 01638.00 | Patch C 01836.00 |
| Patch C 01389.16 | Patch C 01474.03 | Patch C 01642.00 | Patch C 01847.00 |
| Patch C 01392.00 | Patch C 01475.00 | Patch C 01651.00 | Patch C 01849.00 |
| Patch C 01393.00 | Patch C 01480.00 | Patch C 01654.00 | Patch C 01850.00 |
| Patch C 01395.00 | Patch C 01488.00 | Patch C 01655.00 | Patch C 01852.00 |
| Patch C 01397.00 | Patch C 01497.00 | Patch C 01656.00 | Patch C 01853.00 |
| Patch C 01398.00 | Patch C 01500.00 | Patch C 01656.01 | Patch C 01854.00 |
| Patch C 01399.00 | Patch C 01504.00 | Patch C 01656.02 | Patch C 01855.00 |
| Patch C 01399.01 | Patch C 01505.00 | Patch C 01656.03 | Patch C 01856.00 |
| Patch C 01401.00 | Patch C 01507.00 | Patch C 01656.04 | Patch C 01859.00 |
| Patch C 01404.00 | Patch C 01509.02 | Patch C 01657.00 | Patch C 01865.00 |
| Patch C 01406.00 | Patch C 01510.00 | Patch C 01657.01 | Patch C 01865.03 |
| Patch C 01406.01 | Patch C 01511.00 | Patch C 01657.02 | Patch C 01866.00 |
| Patch C 01406.02 | Patch C 01513.00 | Patch C 01657.03 | Patch C 01867.00 |
| Patch C 01406.03 | Patch C 01517.00 | Patch C 01657.04 | Patch C 01868.00 |
| Patch C 01406.10 | Patch C 01519.00 | Patch C 01658.00 | Patch C 01874.00 |
| Patch C 01407.04 | Patch C 01525.00 | Patch C 01662.00 | Patch C 01876.00 |
| Patch C 01411.00 | Patch C 01526.00 | Patch C 01664.00 | Patch C 01880.00 |
| Patch C 01423.00 | Patch C 01526.01 | Patch C 01665.03 | Patch C 01882.00 |
| Patch C 01425.00 | Patch C 01528.00 | Patch C 01666.02 | Patch C 01884.00 |
| Patch C 01429.00 | Patch C 01529.00 | Patch C 01673.00 | Patch C 01886.00 |
| Patch C 01430.00 | Patch C 01530.00 | Patch C 01681.00 | Patch C 01887.00 |
| Patch C 01432.00 | Patch C 01533.00 | Patch C 01683.01 | Patch C 01889.00 |
| Patch C 01433.00 | Patch C 01534.00 | Patch C 01695.00 | Patch C 01890.00 |
| Patch C 01436.00 | Patch C 01535.00 | Patch C 01695.01 | Patch C 01907.00 |
| Patch C 01437.00 | Patch C 01536.00 | Patch C 01704.00 | Patch C 01908.00 |
| Patch C 01439.00 | Patch C 01538.00 | Patch C 01707.00 | Patch C 01909.00 |
| Patch C 01440.00 | Patch C 01541.00 | Patch C 01711.00 | Patch C 01909.01 |
| Patch C 01441.00 | Patch C 01542.00 | Patch C 01716.00 | Patch C 01911.00 |
| Patch C 01442.00 | Patch C 01543.00 | Patch C 01717.00 | Patch C 01916.00 |
| Patch C 01444.00 | Patch C 01545.00 | Patch C 01718.00 | Patch C 01923.01 |
| Patch C 01445.00 | Patch C 01546.00 | Patch C 01721.00 | Patch C 01924.00 |
| Patch C 01448.00 | Patch C 01551.00 | Patch C 01722.00 | Patch C 01925.00 |
| Patch C 01449.00 | Patch C 01552.00 | Patch C 01727.00 | Patch C 01926.00 |
| Patch C 01450.00 | Patch C 01552.01 | Patch C 01728.00 | Patch C 01926.01 |
| Patch C 01450.01 | Patch C 01563.00 | Patch C 01730.00 | Patch C 01927.00 |
| Patch C 01451.00 | Patch C 01564.00 | Patch C 01731.00 | Patch C 01932.00 |
| Patch C 01454.00 | Patch C 01565.00 | Patch C 01732.00 | Patch C 01932.01 |
| Patch C 01455.00 | Patch C 01569.00 | Patch C 01732.01 | Patch C 01933.00 |
| Patch C 01455.01 | Patch C 01570.00 | Patch C 01740.00 | Patch C 01934.00 |
| Patch C 01456.00 | Patch C 01581.00 | Patch C 01758.00 |                  |
| Patch C 01456.01 | Patch C 01582.00 | Patch C 01779.00 |                  |

## New Superseded TruCluster Server CSPs

The numbers in the following table represents TruCluster Server CSPs that are superseded in this kit.

**Table 5-9 New Superseded TruCluster Server CSPs**

|                  |                  |                  |                  |
|------------------|------------------|------------------|------------------|
| Patch C 00428.00 | Patch C 00429.00 | Patch C 00434.00 | Patch C 00450.00 |
|------------------|------------------|------------------|------------------|

## TruCluster Server CSPs Superseded in Previous Kit

The numbers in the following table represents TruCluster Server CSPs that were superseded in previous V5.1B patch kits.

**Table 5-10 TruCluster CSPs Superseded in Previous V5.1B Kits**

|                  |                  |                  |                  |
|------------------|------------------|------------------|------------------|
| Patch C 00390.00 | Patch C 00405.00 | Patch C 00413.00 | Patch C 00425.00 |
| Patch C 00392.00 | Patch C 00406.00 | Patch C 00416.00 | Patch C 00425.01 |
| Patch C 00392.01 | Patch C 00407.00 | Patch C 00419.00 | Patch C 00426.01 |
| Patch C 00393.00 | Patch C 00408.02 | Patch C 00423.00 | Patch C 00427.00 |
| Patch C 00403.00 | Patch C 00409.00 | Patch C 00424.00 | Patch C 00434.00 |

**Table 5-11 TruCluster CSPs Superseded in Previous V5.1B Kits**

|                  |                  |                  |                  |
|------------------|------------------|------------------|------------------|
| Patch C 00005.00 | Patch C 00083.02 | Patch C 00372.01 | Patch C 00274.02 |
| Patch C 00008.00 | Patch C 00083.03 | Patch C 00376.00 | Patch C 00275.00 |
| Patch C 00009.00 | Patch C 00090.00 | Patch C 00203.00 | Patch C 00278.00 |
| Patch C 00010.00 | Patch C 00091.00 | Patch C 00204.00 | Patch C 00278.01 |
| Patch C 00014.00 | Patch C 00092.00 | Patch C 00204.01 | Patch C 00284.00 |
| Patch C 00014.01 | Patch C 00093.00 | Patch C 00205.01 | Patch C 00286.00 |
| Patch C 00017.00 | Patch C 00095.00 | Patch C 00206.00 | Patch C 00286.01 |
| Patch C 00017.01 | Patch C 00096.00 | Patch C 00206.01 | Patch C 00287.00 |
| Patch C 00017.02 | Patch C 00097.00 | Patch C 00207.03 | Patch C 00288.00 |
| Patch C 00020.00 | Patch C 00098.00 | Patch C 00208.00 | Patch C 00290.00 |
| Patch C 00020.01 | Patch C 00099.00 | Patch C 00209.00 | Patch C 00290.01 |
| Patch C 00022.00 | Patch C 00100.00 | Patch C 00210.00 | Patch C 00291.00 |
| Patch C 00024.00 | Patch C 00101.00 | Patch C 00211.00 | Patch C 00292.00 |
| Patch C 00027.00 | Patch C 00102.00 | Patch C 00213.00 | Patch C 00293.02 |
| Patch C 00027.01 | Patch C 00103.00 | Patch C 00215.00 | Patch C 00295.00 |
| Patch C 00029.00 | Patch C 00106.00 | Patch C 00218.00 | Patch C 00296.00 |
| Patch C 00029.01 | Patch C 00109.00 | Patch C 00219.00 | Patch C 00297.00 |
| Patch C 00030.02 | Patch C 00113.00 | Patch C 00219.01 | Patch C 00297.01 |
| Patch C 00031.00 | Patch C 00115.00 | Patch C 00219.02 | Patch C 00300.00 |
| Patch C 00032.00 | Patch C 00116.00 | Patch C 00221.00 | Patch C 00307.00 |
| Patch C 00033.00 | Patch C 00117.00 | Patch C 00224.00 | Patch C 00307.01 |
| Patch C 00034.00 | Patch C 00118.00 | Patch C 00225.01 | Patch C 00316.00 |
| Patch C 00036.00 | Patch C 00119.00 | Patch C 00235.00 | Patch C 00316.01 |
| Patch C 00037.00 | Patch C 00121.00 | Patch C 00235.01 | Patch C 00316.02 |
| Patch C 00037.01 | Patch C 00124.00 | Patch C 00227.00 | Patch C 00317.00 |
| Patch C 00039.00 | Patch C 00125.00 | Patch C 00227.01 | Patch C 00317.03 |
| Patch C 00040.00 | Patch C 00127.00 | Patch C 00227.02 | Patch C 00319.00 |
| Patch C 00041.00 | Patch C 00129.00 | Patch C 00229.00 | Patch C 00320.00 |
| Patch C 00042.00 | Patch C 00132.00 | Patch C 00230.00 | Patch C 00323.00 |
| Patch C 00043.00 | Patch C 00142.01 | Patch C 00230.01 | Patch C 00323.01 |
| Patch C 00045.00 | Patch C 00144.00 | Patch C 00234.00 | Patch C 00327.00 |
| Patch C 00046.00 | Patch C 00146.00 | Patch C 00241.00 | Patch C 00331.00 |
| Patch C 00047.00 | Patch C 00151.00 | Patch C 00243.00 | Patch C 00332.00 |
| Patch C 00048.00 | Patch C 00155.00 | Patch C 00247.00 | Patch C 00335.00 |
| Patch C 00050.00 | Patch C 00158.00 | Patch C 00247.01 | Patch C 00336.00 |
| Patch C 00052.00 | Patch C 00159.00 | Patch C 00247.02 | Patch C 00337.00 |
| Patch C 00052.01 | Patch C 00159.01 | Patch C 00247.03 | Patch C 00338.00 |
| Patch C 00052.02 | Patch C 00159.02 | Patch C 00248.00 | Patch C 00340.00 |
| Patch C 00053.00 | Patch C 00159.03 | Patch C 00249.00 | Patch C 00341.02 |
| Patch C 00053.01 | Patch C 00159.04 | Patch C 00249.01 | Patch C 00343.00 |
| Patch C 00054.00 | Patch C 00159.05 | Patch C 00250.00 | Patch C 00344.00 |
| Patch C 00055.00 | Patch C 00159.06 | Patch C 00250.01 | Patch C 00347.00 |
| Patch C 00056.03 | Patch C 00163.00 | Patch C 00252.00 | Patch C 00348.00 |
| Patch C 00059.00 | Patch C 00165.04 | Patch C 00256.00 | Patch C 00349.00 |
| Patch C 00060.00 | Patch C 00168.00 | Patch C 00258.00 | Patch C 00349.01 |
| Patch C 00061.00 | Patch C 00173.00 | Patch C 00264.00 | Patch C 00349.02 |
| Patch C 00062.00 | Patch C 00174.00 | Patch C 00264.01 | Patch C 00350.00 |
| Patch C 00065.00 | Patch C 00174.01 | Patch C 00266.00 | Patch C 00352.00 |
| Patch C 00066.00 | Patch C 00175.03 | Patch C 00266.01 | Patch C 00352.01 |
| Patch C 00066.02 | Patch C 00176.00 | Patch C 00266.02 | Patch C 00353.00 |
| Patch C 00066.03 | Patch C 00177.01 | Patch C 00266.03 | Patch C 00355.00 |
| Patch C 00066.04 | Patch C 00179.00 | Patch C 00266.04 | Patch C 00358.00 |
| Patch C 00068.00 | Patch C 00181.00 | Patch C 00266.05 | Patch C 00359.00 |
| Patch C 00069.00 | Patch C 00181.01 | Patch C 00270.08 | Patch C 00361.00 |
| Patch C 00069.01 | Patch C 00182.00 | Patch C 00271.00 | Patch C 00364.00 |
| Patch C 00071.00 | Patch C 00188.00 | Patch C 00380.00 | Patch C 00365.00 |
| Patch C 00072.00 | Patch C 00190.00 | Patch C 00384.00 | Patch C 00366.00 |
| Patch C 00073.00 | Patch C 00193.00 | Patch C 00379.00 | Patch C 00367.00 |
| Patch C 00076.00 | Patch C 00193.01 | Patch C 00383.00 | Patch C 00368.00 |
| Patch C 00078.00 | Patch C 00193.02 | Patch C 00386.00 | Patch C 00369.00 |
| Patch C 00079.00 | Patch C 00194.00 | Patch C 00392.00 | Patch C 00369.01 |
| Patch C 00081.00 | Patch C 00194.01 | Patch C 00392.01 | Patch C 00372.00 |
| Patch C 00082.00 | Patch C 00197.01 | Patch C 00387.00 |                  |
| Patch C 00083.00 | Patch C 00199.01 | Patch C 00274.00 |                  |
| Patch C 00083.01 | Patch C 00200.00 | Patch C 00274.01 |                  |



# A Setting Up an Enhanced Distance Cluster

An Enhanced Distance Cluster allows a cluster to be extended between two sites up to 100 km apart to assist recovery in the event of a disaster. An Enhanced Distance Cluster provides basic high availability services in the event of the loss of a single component. It is important to note that an Enhanced Distance Cluster is not designed to handle nor is it capable of handling simultaneous cascading failures and therefore can not provide a fully disaster tolerant solution. The following topics are discussed:

- The section “Enhanced Distance Cluster Configuration Requirements” lists configurations, features, and functions that must be present in order to have a supported Enhanced Distance Cluster configuration.
- The section “Enhanced Distance Cluster Restrictions” (page 192) lists configurations, features, and functions that cannot be present in order to have a supported Enhanced Distance Cluster configuration.
- The section “Enhanced Distance Cluster Recommendations” (page 193) lists configurations, features, functions that should be present, but are not required to have a working Enhanced Distance Cluster.

## Enhanced Distance Cluster Configuration Requirements

This section provides information on the hardware and configurations supported for Enhanced Distance Clusters. Configurations that deviate from these configurations will require a custom support statement from Hewlett Packard in order to be supported. See the TruCluster Server QuickSpecs for information on hardware components supported by the TruCluster Server product.

The following table lists supported hardware.

**Table A-1 Hardware Configuration**

|                             |                                                                                                                                                   |
|-----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------|
| Number of Nodes             | 2 to 4 nodes                                                                                                                                      |
| CPU Type                    | Any node supported in a TruCluster configuration.                                                                                                 |
| Fibre Channel Adapter (HBA) | Any Fibre Channel HBA supported by TruCluster. (Parallel SCSI is not supported for any storage that can potentially fail over between locations.) |
| Interconnect                | Gigabit Ethernet LAN with up to three switches. Cumulative distance must not exceed 100 km.                                                       |
| Storage                     | TruCluster Fibre Channel connected storage (such as HP StorageWorks XP and HP StorageWorks EVA).                                                  |

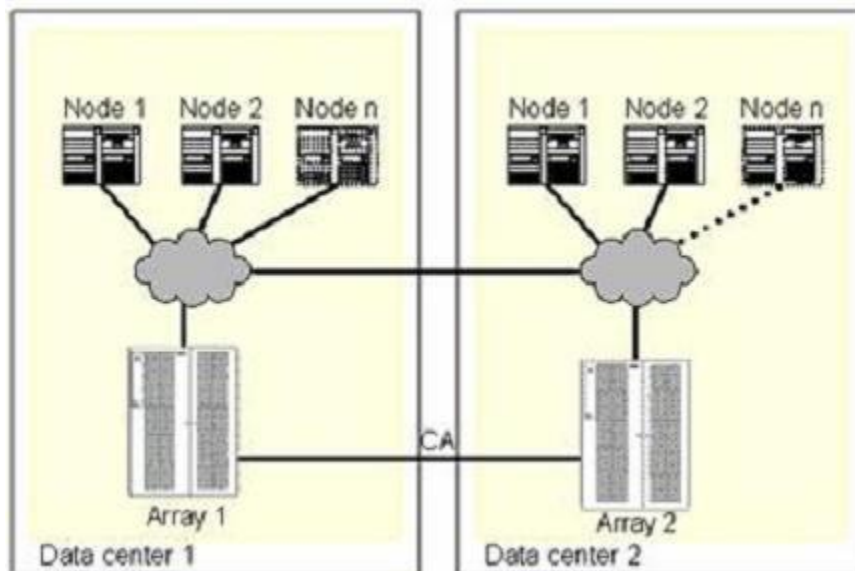
The following list describes configuration requirements for the cluster:

- A two- to four-node cluster with up to three nodes at one site and the other nodes at a different site.
- The cluster must be configured as a LAN cluster.
- A least one shared storage array must be present at each site.
- The cluster root (/), /usr, and /var file systems must be located within the same site.
- All SAN-attached storage must be shared and directly accessible — not over the cluster inter-site connection but via the SAN — from all nodes at both sites.
- The storage must be configured with remote data replication software (such as XP Continuous Access). Data replication is required in order to provide the ability to boot the site that does not contain the cluster file systems following a disaster event.
- A reboot of all nodes at the surviving site is required following any disaster event that requires activation of the secondary replicated volumes. You will need to shut down the system, reconfigure the storage as necessary (perform an XP takeover, for example), and reboot the system. The expected quorum votes and other parameters may need to be modified in order to successfully boot the system.

- If a site disaster occurs that involves multiple failures, high availability will be lost. Therefore, there needs to be procedures in place for the manual rebooting of the surviving site. The surviving site will work as a normal cluster with minimal or no data loss.
- A single, combined span of up to 100 km using three switches and two segments of 50 km.
- The configuration must have at least one physical subnet to which all cluster members and the default cluster alias belong.
- The cluster must have an extended, dedicated cluster interconnect to which all cluster members are connected to serve as a private communication channel between cluster members. The interconnect must be shielded from any traffic that is not a part of the cluster communication according to the requirements for the LAN based cluster interconnect in the TruCluster *Cluster Hardware Configuration* manual.

Figure A-1 provides an example of an Enhanced Distance Cluster cluster interconnect configuration. The nodes at Data Center 1 are connected to a switch that is connected to an intermediate switch using a fiber link of up to 50 km. From the intermediate switch, another fiber link of up to 50 km connects to a third switch, to which the remaining two nodes at Data Center 2 are attached, thereby, establishing an overall distance between the sites of up to 100 km.

**Figure A-1 Enhanced Distance Cluster Configuration**



## Enhanced Distance Cluster Restrictions

Although an Enhanced Distance Cluster provides basic high availability between remote sites at greater distances, there are additional restrictions necessary to reduce the volume of data and its latency. Issues and restrictions are as follows:

- The use of the Logical Storage Manager (LSM) to mirror user data or system disk/data in an Enhanced Distance Cluster configuration is prohibited. LSM can only be used for striping data in an Enhanced Distance Cluster.
- Only single-instance applications are supported in an Enhanced Distance Cluster configuration. Multi-instance applications that use the cluster interconnect for communication between multiple instances, such as ORACLE RAC, are not recommended in a Enhanced Distance Cluster configuration. This is due to increased latency on the cluster interconnect.



For example, ORACLE RAC uses a cluster interconnect for its own lock manager traffic and an increase in latency could lead to unpredictable results.

Additionally, multi-instance applications can increase the burden on the cluster interconnect and affect inter-member cluster communication so should follow the configuration recommendations in the next section.

- An Enhanced Distance Cluster should not be used for the following workloads due to their I/O characteristics and directory locations. (Mail and print usually keep their data stores in the system directories):
  - NFS server
  - Mail server
  - Print server

## Enhanced Distance Cluster Recommendations

Although Enhanced Distance Cluster configurations provide access from both sites to all shared data storage in the system, the volume of traffic to off-site storage, either through the storage or cluster inter-site connections, should be minimized to maintain performance levels. Therefore, you should be aware of the following recommendations are applicable. Not following these recommendations can have detrimental effect on system performance and availability.

The main purpose of most of these recommendations is to design a configuration that minimizes I/O latency and contention issues within both the cluster and the SAN storage inter-site connections:

- The configuration of an application and its data storage should be located within the same site. Doing otherwise can cause significant performance degradation to the application's runtime performance.
- Each application should have its associated data in separate file systems, or AdvFS domains, in order to locate the CFS server for these file systems with the application.
- Applications should not use any storage/directory within the system tree (for example, root, /usr, /tmp, /var, /var/tmp, and so on). Instead they should use dedicated file systems that can be co-located with the application. If the application requires a directory in the system tree, a symbolic link should be used to redirect these directories to a separate storage/file system.
- The use of remote I/O via DRD to serve devices is strongly discouraged due to the increased traffic on the cluster interconnect possibly resulting in unacceptable latencies.
- Although the use of LSM for striping is supported, it is not recommended for use.
- Although data replication is supported between HSG80 arrays, their use as shared storage within an Enhanced Distance Cluster it is not recommended. This is due to the inter-site distance limitations of the HSG80's remote data replication, which would significantly impact the Enhanced Distance Cluster configuration.
- The use of a cluster alias for network-client access is discouraged due to the potential for additional traffic on the cluster interconnect if the alias router is not located with the application.

If a cluster alias is used, each application must have its own cluster alias. Each application should have its cluster alias router located on the node where the application is executing. This can be achieved by using the `cluamgr` command and adjusting the values of selection weight and routing priority. During a failover of the application, the cluster alias also should be relocated to the same node.



---

## B Prior Patch Installation Changes

Beginning with Version 5.1B-2, we changed the way Tru64 UNIX patch kits are installed and removed, and Version 5.1B-3 introduced additional changes. The following sections describe these changes.

See the [Patch Kit Installation Instructions](#) manual for complete information about installing, removing, and managing Tru64 UNIX patch kits.

### Changes Made in Version 5.1B-3

If you did not install Version 5.1B-3 but have installed earlier kits, you will see the following differences in the kitting process:

- Before you can install this kit you must accept the conditions included in the license that the `dupatch` utility displays. You can read this license in [Appendix C “Component Licensing”](#).
- You can delete the patch kit by kit name rather than by specifying individual patches. With the introduction of this feature, you can easily delete patches interactively using the `dupatch` utility or from the `dupatch` command line. This feature also works on pre-Version 5.1B-6 kits.
- You can delete patches in multi-user mode.
- You can force the installation of the patch kit even if file conflicts exist. This feature is an extension of the `dupatch` baselining feature.
- A new command-line option, Patch Level, provides a single command that provides a full description of the patch kits, CSPs, and ERPs installed on your system.

### Changes Made in Version 5.1B-2

If you did not install versions 5.1B-2 or 5.1B-3 but have installed earlier kits, you will see the following differences in the kitting process:

- All or none installation

When you install an inclusive patch kit, you must install all patches; you can no longer select specific patches to install. By making the installation of all patches mandatory, you can patch with greater confidence that the process will be problem free.

Before a patch kit is released, it is tested on many types of systems and system configurations. This testing continues until we are assured that the patches perform the tasks they were designed for and do not introduce new problems. It is not possible to achieve this type of testing on every possible combination of individually selected patches.

- Substantially reduced installation time

The installation process for inclusive patch kits can reduce the time it takes to install the patches by as much as half from what you are used to. For large, clustered systems, the difference can be several hours faster.

- Fewer patches displayed

Because of the way these new patch kits are designed, you will see many fewer patches listed by `dupatch` during the installation process. For example, a partial listing you see will be similar to the following:

```
- Tru64_UNIX_V5.1B
/ Security Related Patches:
 * Patch 27001.00 - SP04 OSFACCT540

 * Patch 27002.00 - SP04 OSFADVFS540 (SSRT2275)

 * Patch 27003.00 - SP04 OSFADVFSBIN540
```

In the old-style patch kits, these three patches might have consisted of perhaps 20 individual patches being displayed. The difference is not in the content of the kits, but rather in the

way the patches are packaged and installed. In this example, the SPO4 identifies the patch as belonging to Version 5.1B-2 (Patch Kit 4), the OSF . . . 540 identifies the subset the patch is included in, and the SSRT2275 indicates a type of security patch.

- All or none patch removal  
As with the installation process, if you want to remove a patch, you must remove all of them. That is, you can no longer select individual patches for removal.
- Patches for Worldwide Language Support (WLS) subset  
The inclusive patch kits deliver patches that may be required for the WLS subset. As with the TruCluster Server patches, the WLS patches will only be installed if you have the WLS subset installed. See Chapter 4 “Worldwide Language Support Patches” for information about WLS subset patches.

---

# C Component Licensing

This appendix provides the licenses for software components included in this kit.

## HP Tru64 UNIX Version 5.1B-6 Consolidated Patch Kit

ATTENTION: USE OF THE SOFTWARE IS SUBJECT TO THE HP Tru64 UNIX Version 5.1B-6 Consolidated Patch Kit SOFTWARE LICENSE AGREEMENT BELOW. YOU MUST REVIEW THE AGREEMENT AND EITHER ACCEPT OR NOT ACCEPT THE AGREEMENT. IF YOU DO NOT ACCEPT THE AGREEMENT, YOU MAY NOT USE THE SOFTWARE.

HP IS WILLING TO LICENSE THE SOFTWARE IDENTIFIED ABOVE (THE "SOFTWARE") INCLUDING, WITHOUT LIMITATION, ANY DOCUMENTATION AND ANY ANCILLARY SOFTWARE BUNDLED WITH OR EMBEDDED IN THE SOFTWARE (COLLECTIVELY, THE "ANCILLARY SOFTWARE"), TO YOU SUBJECT TO THE SOFTWARE LICENSE TERMS AND THE APPLICABLE "LIMITED WARRANTY STATEMENT" PROVIDED WITH THE SOFTWARE.

HP HAS IDENTIFIED ANCILLARY SOFTWARE BY EITHER NOTING THE RELEVANT PARTY'S OWNERSHIP WITHIN EACH ANCILLARY SOFTWARE PROGRAM FILE AND/OR BY PROVIDING INFORMATION IN THE "ReleaseNotes.pdf " FILE THAT IS PROVIDED AS PART OF THE SOFTWARE. YOUR USE OF ANY ANCILLARY SOFTWARE SHALL BE GOVERNED BY THAT PARTY'S LICENSE AGREEMENT ("ANCILLARY SOFTWARE LICENSE") AND NOT BY THIS AGREEMENT. THE LICENSES FOR THE ANCILLARY SOFTWARE ARE INCLUDED IN SUCH A NCILLARY SOFTWARE AND/OR SET FORTH IN THE "ReleaseNotes.pdf " FILE THAT IS PROVIDED AS PART OF THE SOFTWARE. IF YOU CHOOSE TO ACCEPT THIS AGREEMENT WITHOUT REVIEWING SUCH ANCILLARY SOFTWARE LICENSES OR OTHER TERMS, YOU WILL BE DEEMED TO HAVE ACCEPTED SUCH ANCILLARY LICENSES OR OTHER TERMS.

BY USING THE SOFTWARE YOU INDICATE YOUR ACCEPTANCE OF THE SOFTWARE LICENSE TERMS AND THE APPLICABLE WARRANTY STATEMENTS CONTAINED IN OR OTHERWISE ACCOMPANYING THE SOFTWARE, AS WELL AS THE TERMS AND CONDITIONS ACCOMPANYING THE ANCILLARY SOFTWARE SET FORTH IN THE ReleaseNotes.pdf FILE. IF YOU DO NOT ACCEPT SUCH SOFTWARE LICENSE AND WARRANTY STATEMENT TERMS, AS WELL AS THE TERMS AND CONDITIONS ACCOMPANYING THE ANCILLARY SOFTWARE OR SET FORTH THE ReleaseNotes.pdf FILE, THEN YOU ARE NOT GRANTED A LICENSE TO THE SOFTWARE, YOU ARE NOT AUTHORIZED TO USE THE SOFTWARE, AND YOU MAY NOT DOWNLOAD THE SOFTWARE.

NOTICE: This software is licensed, not sold, if you do not agree to these terms, you may not use the software.

### 1. DEFINITIONS

- a) "Use" means downloading, installing, storing, executing, or displaying Software on a Device.
- b) "Products" means Software, documentation, accessories, supplies, and upgrades that are determined by HP to be available from HP upon receipt of Customer's order.
- c) "Software License" means the Software license grant and general license terms set forth herein.

### 2. LICENSE GRANT

- a) Provided that Customer has a valid software license to use HP Tru64 UNIX Version 5.1B-6 (as described below), HP grants Customer a world-wide, non-exclusive license to Use the object code version of the Software so long as it is used solely in conformance with:

1) The terms set forth herein; and

2) HP's third party suppliers' terms that accompany the Software. In the event of a conflict, the third party suppliers' terms that accompany the Software will take precedence over the terms set forth herein. The terms applicable to this transaction in total are referred to as the "Agreement". A valid software license to use HP Tru64 UNIX Version 5.1B-6 can be purchased separately or through the ownership of a valid software update license which can be purchased separately or via an update support agreement which includes license to use, license subscription, or rights to use new versions of software.

b) All Software Licenses will be perpetual unless terminated or transferred in accordance with Section 3. g).

c) If Customer is an HP authorized reseller, Customer may sublicense the Software to an end-user for its Use or (if applicable) sublicense the Software to an HP authorized reseller for subsequent distribution to an end-user for its Use. These sublicenses must incorporate the terms of this Software License in a written sublicense agreement, which will be made available to HP upon request. If Customer is not an HP authorized reseller, Customer may not sublicense the Software unless otherwise agreed to by HP in writing.

e) HP, or its designee(s), shall, during regular business hours at Customer's offices and in such a manner that does not interfere with Customer's normal business activities, have the right to inspect and audit, or have a third party perform an inspection and audit, the number of copies of Software Used or distributed by Customer, the computers on which the Software, if any, is installed and the number of users Using any such Software. If any audit discloses underpayments of five percent (5%) or more of the amount of license fees Customer should have actually paid to HP for the valid software license to use HP Tru64 UNIX Version 5.1B-6 entitling Customer to Use the Software, Customer shall bear all of the costs of the audit. HP's audit rights shall not terminate or expire until three (3) years after termination or expiration of this Agreement.

### 3. GENERAL LICENSE TERMS

a) Software is owned and copyrighted by HP or by third party suppliers. Customer's Software License confers no title or ownership and is not a sale of any rights in the Software. Third party suppliers are intended beneficiaries under this Agreement and may protect their rights in the Software directly against Customer in the event of any infringement.

b) Unless otherwise permitted in writing by HP, Customer may only make copies or adaptations of the Software for archival purposes or when copying or adaptation is an essential step in the authorized Use of the Software on a backup Device, provided that copies and adaptations are used in no other manner and provided further that the Use on the backup Device is discontinued when the original or replacement Device becomes operable.

c) Customer must reproduce all copyright notices and other proprietary legends in or on the original Software on all permitted copies or adaptations. Customer may not copy the Software onto any public or distributed network.

d) Subject to Section 3.b), above, Bundled Software provided to Customer may only be used when operating the Device in which the Bundled Software was installed by HP in configurations as sold or subsequently upgraded by HP.

e) Updates, upgrades or other enhancements are only available under HP Support agreements. HP reserves the right to require additional licenses and fees for Use of the Software on Devices other than the Devices in which the Bundled Software were installed by HP.

f) Customer will not modify, disassemble or decompile the Software without HP's prior written consent. Where Customer has other rights under statute, Customer will provide HP with reasonably detailed information regarding any intended disassembly or decompilation. Customer will not decrypt the Software unless necessary for legitimate use of the Software.

g) Customer's Software License is transferable only upon HP's prior written authorization and payment to HP of any applicable fee(s). Upon transfer of the Software License, Customer will immediately deliver all copies of the Software to the transferee. The transferee must agree in writing to the terms of Customer's Software License. All Software License terms will be binding on involuntary transferees, notice of which is hereby given. Any transfer of the Software by the Customer in accordance with the provisions of this Section 3.g) shall not relieve Customer of any liability for the performance of Customer's obligations under this Agreement. Customer's right to Use the Software License will automatically terminate upon transfer.

h) HP may terminate Customer's or any transferee's or sublicensee's Software License upon notice for failure to comply with any applicable Agreement terms. Immediately upon termination, the Software and all copies of the Software will be destroyed or returned to HP. Copies of the Software that are merged into adaptations, except for individual pieces of data in Customer's or transferee's or sublicensee's database, will be removed and destroyed or returned to HP. With HP's written consent, one copy of the Software may be retained subsequent to termination for archival purposes. Customer may terminate this Agreement by returning the Software to HP.

i) In the following provision regarding Software Licenses to the U.S. Government, the term "Customer" means HP's direct purchaser, any entity or individual sublicensing the Software, and the end-user.

1) If Software is licensed for use in the performance of a U.S government prime contract or subcontract, Customer agrees that Software has been developed entirely at private expense. Customer agrees that Software, and any derivatives or modifications, is adequately marked when the Restricted Rights Legend below is affixed to the Software or to its storage media and is perceptible directly or with the aid of a machine or device. Customer agrees to conspicuously put the following legend on the Software media with Customer's name and address added below the notice:

RESTRICTED RIGHTS LEGEND

Use, duplication or disclosure is subject to HP standard commercial license terms and for non-DOD Departments and Agencies of the U.S. Government, the restrictions as set forth in FAR 52.227-19(c) (1-2) (Jun 1987).

Hewlett-Packard Company  
3000 Hanover Street  
Palo Alto, CA 94304 U.S.A.

Copyright (c) 2005 Hewlett-Packard Company. All Rights Reserved

2) Customer further agrees that Software is delivered and licensed as "Commercial computer software" as defined in

DFARS 252.227-7014 (Jun 1995) or as a "commercial item" as defined in FAR 2.101(a), or as "Restricted computer software" as defined in FAR 52.227-19 (Jun 1987) (or any equivalent agency regulation or contract clause), whichever is applicable. The Customer agrees that it has only those rights provided for such Software by the applicable FAR or DFARS clause or the HP standard software agreement for the product involved.

#### 4. GENERAL

- a) Customer may not assign or transfer this Agreement or any rights or obligations hereunder without prior written consent of HP. Any such attempted assignment or transfer will be null and void. HP may terminate this Agreement in the event of any such attempted assignment or transfer.
  
- b) Customer agrees and certifies that the Software, HP licensed Products, technology or technical data obtained hereunder, will not be exported, re-exported or imported except as authorized and permitted by the laws and regulations for obtaining required export and import authorizations. HP may terminate this Agreement if Customer is in violation of any applicable laws or regulations. Customer agrees to indemnify and hold HP and its suppliers harmless from any claims arising from any breach of this Section by Customer or any of its directors, officers, employees, subdistributors, customers or any third party to whom Customer provides access to the Software.
  
- c) Disputes arising in connection with this Agreement will be governed by the laws of the country and locality in which HP accepts the order. For orders accepted in the United States, disputes will be governed by the laws of the state of California.
  
- d) If any term or provision herein is determined to be illegal or unenforceable, the validity or enforceability of the remainder of the terms or provisions herein will remain in full force and effect. Provisions herein which by their nature extend beyond the termination of any license of Software will remain in effect until fulfilled.
  
- e) These HP Software License Terms supersede any previous communications, representations or agreements between the parties, whether oral or written, regarding transactions hereunder. Customer's additional or different terms and conditions will not apply. These HP Software License Terms may not be changed except by an amendment signed by an authorized representative of each party.
  
- f) The waiver or failure of either party to exercise in any respect any right provided for in the Agreement shall not be deemed a waiver of any further right under this Agreement.

#### LIMITED WARRANTY STATEMENT

HP SOFTWARE PRODUCT DURATION OF LIMITED WARRANTY  
90 Days {MATCH TO HP WARRANTY CODE}

Scope. The limited warranty is limited to the HP owned software portion of the HP software product ("Software"). The warranty for any other software portion of the HP software product ("Ancillary Software"), if any, shall be governed by the warranty terms provided with the Ancillary Software and to the extent allowed by local law HP is providing the Ancillary Software "AS-IS" WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, WHETHER ORAL OR WRITTEN, EXPRESS OR IMPLIED, AND HP SPECIFICALLY DISCLAIMS ANY IMPLIED WARRANTIES OR CONDITIONS OF MERCHANTABILITY, SATISFACTORY QUALITY, NON-INFRINGEMENT, TITLE, ACCURACY OF



INFORMATIONAL CONTENT, AND FITNESS FOR A PARTICULAR PURPOSE WITH RESPECT TO THE ANCILLARY SOFTWARE; THE ENTIRE RISK AS TO THE RESULTS AND PERFORMANCE OF THE ANCILLARY SOFTWARE IS ASSUMED BY YOU.

Software Limited Warranty. HP warrants to you that the Software will not fail to execute its programming instructions after the date of purchase, for the period specified above, due to defects in material and workmanship when properly installed and used.

If HP receives notice of such defects during the warranty period, HP will replace Software that does not execute its programming instructions due to such defects.

HP does not warrant that the operation of Software will be uninterrupted or error free.

If HP is unable, within a reasonable time, to repair or replace any product to a condition as warranted, you will be entitled to a refund of the purchase price upon prompt return of the product.

Exclusions. This limited warranty does not apply to defects resulting from (a) improper or inadequate maintenance, (b) unauthorized modification or misuse, or (c) operation outside of the published environmental specifications for the product or otherwise in an unclean environment.

Disclaimer. TO THE EXTENT ALLOWED BY LOCAL LAW, THE ABOVE WARRANTIES ARE EXCLUSIVE AND NO OTHER WARRANTY OR CONDITION, WHETHER WRITTEN OR ORAL, IS EXPRESSED OR IMPLIED. HP SPECIFICALLY DISCLAIMS ANY IMPLIED WARRANTIES OR CONDITIONS OF MERCHANTABILITY, SATISFACTORY QUALITY, NON-INFRINGEMENT, TITLE, ACCURACY OF INFORMATIONAL CONTENT, AND FITNESS FOR A PARTICULAR PURPOSE. HP DOES NOT WARRANT THAT THE FUNCTIONS CONTAINED IN THE SOFTWARE WILL MEET YOUR REQUIREMENTS. THE ENTIRE RISK AS TO THE RESULTS AND PERFORMANCE OF THE SOFTWARE IS ASSUMED BY YOU. NO ORAL OR WRITTEN INFORMATION OR ADVICE GIVEN BY HP OR HP'S AUTHORIZED REPRESENTATIVES SHALL CREATE A WARRANTY OR IN ANY WAY INCREASE THE SCOPE OF THIS WARRANTY. Some jurisdictions do not allow limitations on the duration of an implied warranty, so the above limitation or exclusion might not apply to you to the extent prohibited by such local laws. This warranty gives you specific legal rights and you might also have other rights that vary from country to country, state to state, or province to province.

Limitation of Liability. TO THE EXTENT ALLOWED BY LOCAL LAW, THE REMEDIES IN THIS WARRANTY STATEMENT ARE YOUR SOLE AND EXCLUSIVE REMEDIES. EXCEPT TO THE EXTENT PROHIBITED BY LOCAL LAW, IN NO EVENT WILL HP OR ITS SUBSIDIARIES, AFFILIATES, DIRECTORS, OFFICERS, EMPLOYEES, AGENTS OR SUPPLIERS BE LIABLE FOR SPECIAL, INCIDENTAL, CONSEQUENTIAL, PUNITIVE OR OTHER DAMAGES (INCLUDING LOST PROFIT, LOST DATA, OR DOWNTIME COSTS), ARISING OUT OF THE USE, INABILITY TO USE, OR THE RESULTS OF USE OF THE SOFTWARE, WHETHER BASED IN WARRANTY, CONTRACT, TORT OR OTHER LEGAL THEORY, AND WHETHER OR NOT HP WAS ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. THE SOFTWARE IS NOT SPECIFICALLY DESIGNED, MANUFACTURED OR INTENDED FOR USE IN THE PLANNING, CONSTRUCTION, MAINTENANCE, OR DIRECT OPERATION OF A NUCLEAR FACILITY, AIRCRAFT NAVIGATION OR AIRCRAFT COMMUNICATION SYSTEMS, AIR TRAFFIC CONTROL, DIRECT LIFE SUPPORT MACHINES, OR WEAPONS SYSTEMS. CUSTOMER IS SOLELY LIABLE IF THE SOFTWARE IS USED FOR THESE APPLICATIONS. CUSTOMER WILL INDEMNIFY AND HOLD HP HARMLESS FROM ALL LOSS, DAMAGE, EXPENSE OR LIABILITY IN CONNECTION WITH SUCH USE. IN ANY CASE, HP'S ENTIRE LIABILITY UNDER ANY PROVISION

OF THIS AGREEMENT SHALL BE LIMITED TO THE GREATER OF THE AMOUNT ACTUALLY PAID BY YOU FOR THE SOFTWARE OR U.S.\$5.00. Your use of the Software is entirely at your own risk. Some jurisdictions do not allow the exclusion or limitation of liability for incidental or consequential damages, so the above limitation may not apply to you to the extent prohibited by such local laws.

Note. EXCEPT TO THE EXTENT ALLOWED BY LOCAL LAW, THESE WARRANTY TERMS DO NOT EXCLUDE, RESTRICT OR MODIFY, AND ARE IN ADDITION TO, THE MANDATORY STATUTORY RIGHTS APPLICABLE TO THE LICENSE OF THE SOFTWARE TO YOU; PROVIDED, HOWEVER, THAT THE CONVENTION ON CONTRACTS FOR THE INTERNATIONAL SALE OF GOODS IS SPECIFICALLY DISCLAIMED AND SHALL NOT GOVERN OR APPLY TO THE SOFTWARE PROVIDED IN CONNECTION WITH THIS WARRANTY STATEMENT.

BY PRESSING THE CHARACTER "y" AND USING THE SOFTWARE YOU INDICATE YOUR ACCEPTANCE OF THE AGREEMENT. IF YOU PRESS THE CHARACTER "n" AND DO NOT ACCEPT THE AGREEMENT, THEN YOU ARE NOT GRANTED A LICENSE TO THE SOFTWARE, YOU ARE NOT AUTHORIZED TO USE THE SOFTWARE, AND YOU MAY NOT DOWNLOAD THE SOFTWARE.

## Apache License

Apache License  
Version 2.0, January 2004  
<http://www.apache.org/licenses/>

### TERMS AND CONDITIONS FOR USE, REPRODUCTION, AND DISTRIBUTION

#### 1. Definitions.

"License" shall mean the terms and conditions for use, reproduction, and distribution as defined by Sections 1 through 9 of this document.

"Licensor" shall mean the copyright owner or entity authorized by the copyright owner that is granting the License.

"Legal Entity" shall mean the union of the acting entity and all other entities that control, are controlled by, or are under common control with that entity. For the purposes of this definition, "control" means (i) the power, direct or indirect, to cause the direction or management of such entity, whether by contract or otherwise, or (ii) ownership of fifty percent (50%) or more of the outstanding shares, or (iii) beneficial ownership of such entity.

"You" (or "Your") shall mean an individual or Legal Entity exercising permissions granted by this License.

"Source" form shall mean the preferred form for making modifications, including but not limited to software source code, documentation source, and configuration files.

"Object" form shall mean any form resulting from mechanical transformation or translation of a Source form, including but not limited to compiled object code, generated documentation, and conversions to other media types.

"Work" shall mean the work of authorship, whether in Source or Object form, made available under the License, as indicated by a copyright notice that is included in or attached to the work (an example is provided in the Appendix below).

"Derivative Works" shall mean any work, whether in Source or Object form, that is based on (or derived from) the Work and for which the editorial revisions, annotations, elaborations, or other modifications represent, as a whole, an original work of authorship. For the purposes of this License, Derivative Works shall not include works that remain

separable from, or merely link (or bind by name) to the interfaces of, the Work and Derivative Works thereof.

"Contribution" shall mean any work of authorship, including the original version of the Work and any modifications or additions to that Work or Derivative Works thereof, that is intentionally submitted to Licensor for inclusion in the Work by the copyright owner or by an individual or Legal Entity authorized to submit on behalf of the copyright owner. For the purposes of this definition, "submitted" means any form of electronic, verbal, or written communication sent to the Licensor or its representatives, including but not limited to communication on electronic mailing lists, source code control systems, and issue tracking systems that are managed by, or on behalf of, the Licensor for the purpose of discussing and improving the Work, but excluding communication that is conspicuously marked or otherwise designated in writing by the copyright owner as "Not a Contribution."

"Contributor" shall mean Licensor and any individual or Legal Entity on behalf of whom a Contribution has been received by Licensor and subsequently incorporated within the Work.

2. Grant of Copyright License. Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable copyright license to reproduce, prepare Derivative Works of, publicly display, publicly perform, sublicense, and distribute the Work and such Derivative Works in Source or Object form.
3. Grant of Patent License. Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable (except as stated in this section) patent license to make, have made, use, offer to sell, sell, import, and otherwise transfer the Work, where such license applies only to those patent claims licensable by such Contributor that are necessarily infringed by their Contribution(s) alone or by combination of their Contribution(s) with the Work to which such Contribution(s) was submitted. If You institute patent litigation against any entity (including a cross-claim or counterclaim in a lawsuit) alleging that the Work or a Contribution incorporated within the Work constitutes direct or contributory patent infringement, then any patent licenses granted to You under this License for that Work shall terminate as of the date such litigation is filed.
4. Redistribution. You may reproduce and distribute copies of the Work or Derivative Works thereof in any medium, with or without modifications, and in Source or Object form, provided that You meet the following conditions:
  - (a) You must give any other recipients of the Work or Derivative Works a copy of this License; and
  - (b) You must cause any modified files to carry prominent notices stating that You changed the files; and
  - (c) You must retain, in the Source form of any Derivative Works that You distribute, all copyright, patent, trademark, and attribution notices from the Source form of the Work, excluding those notices that do not pertain to any part of the Derivative Works; and
  - (d) If the Work includes a "NOTICE" text file as part of its distribution, then any Derivative Works that You distribute must include a readable copy of the attribution notices contained within such NOTICE file, excluding those notices that do not pertain to any part of the Derivative Works, in at least one of the following places: within a NOTICE text file distributed as part of the Derivative Works; within the Source form or documentation, if provided along with the Derivative Works; or, within a display generated by the Derivative Works, if and wherever such third-party notices normally appear. The contents of the NOTICE file are for informational purposes only and

do not modify the License. You may add Your own attribution notices within Derivative Works that You distribute, alongside or as an addendum to the NOTICE text from the Work, provided that such additional attribution notices cannot be construed as modifying the License.

You may add Your own copyright statement to Your modifications and may provide additional or different license terms and conditions for use, reproduction, or distribution of Your modifications, or for any such Derivative Works as a whole, provided Your use, reproduction, and distribution of the Work otherwise complies with the conditions stated in this License.

5. Submission of Contributions. Unless You explicitly state otherwise, any Contribution intentionally submitted for inclusion in the Work by You to the Licensor shall be under the terms and conditions of this License, without any additional terms or conditions. Notwithstanding the above, nothing herein shall supersede or modify the terms of any separate license agreement you may have executed with Licensor regarding such Contributions.
6. Trademarks. This License does not grant permission to use the trade names, trademarks, service marks, or product names of the Licensor, except as required for reasonable and customary use in describing the origin of the Work and reproducing the content of the NOTICE file.
7. Disclaimer of Warranty. Unless required by applicable law or agreed to in writing, Licensor provides the Work (and each Contributor provides its Contributions) on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied, including, without limitation, any warranties or conditions of TITLE, NON-INFRINGEMENT, MERCHANTABILITY, or FITNESS FOR A PARTICULAR PURPOSE. You are solely responsible for determining the appropriateness of using or redistributing the Work and assume any risks associated with Your exercise of permissions under this License.
8. Limitation of Liability. In no event and under no legal theory, whether in tort (including negligence), contract, or otherwise, unless required by applicable law (such as deliberate and grossly negligent acts) or agreed to in writing, shall any Contributor be liable to You for damages, including any direct, indirect, special, incidental, or consequential damages of any character arising as a result of this License or out of the use or inability to use the Work (including but not limited to damages for loss of goodwill, work stoppage, computer failure or malfunction, or any and all other commercial damages or losses), even if such Contributor has been advised of the possibility of such damages.
9. Accepting Warranty or Additional Liability. While redistributing the Work or Derivative Works thereof, You may choose to offer, and charge a fee for, acceptance of support, warranty, indemnity, or other liability obligations and/or rights consistent with this License. However, in accepting such obligations, You may act only on Your own behalf and on Your sole responsibility, not on behalf of any other Contributor, and only if You agree to indemnify, defend, and hold each Contributor harmless for any liability incurred by, or claims asserted against, such Contributor by reason of your accepting any such warranty or additional liability.

END OF TERMS AND CONDITIONS

APPENDIX: How to apply the Apache License to your work.

To apply the Apache License to your work, attach the following boilerplate notice, with the fields enclosed by brackets "[]" replaced with your own identifying information. (Don't include the brackets!) The text should be enclosed in the appropriate comment syntax for the file format. We also recommend that a file or class name and description of purpose be included on the same "printed page" as the copyright notice for easier identification within third-party archives.

Copyright [yyyy] [name of copyright owner]

Licensed under the Apache License, Version 2.0 (the "License");  
you may not use this file except in compliance with the License.  
You may obtain a copy of the License at

<http://www.apache.org/licenses/LICENSE-2.0>

Unless required by applicable law or agreed to in writing, software  
distributed under the License is distributed on an "AS IS" BASIS,  
WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied.  
See the License for the specific language governing permissions and  
limitations under the License.

## APACHE HTTP SERVER SUBCOMPONENTS

The Apache HTTP Server includes a number of subcomponents with separate copyright notices and license terms. Your use of the source code for these subcomponents is subject to the terms and conditions of the following licenses.

### mod\_mime\_magic component

For the mod\_mime\_magic component:

```
/*
 * mod_mime_magic: MIME type lookup via file magic numbers
 * Copyright (c) 1996-1997 Cisco Systems, Inc.
 *
 * This software was submitted by Cisco Systems to the Apache Group in July
 * 1997. Future revisions and derivatives of this source code must
 * acknowledge Cisco Systems as the original contributor of this module.
 * All other licensing and usage conditions are those of the Apache Group.
 *
 * Some of this code is derived from the free version of the file command
 * originally posted to comp.sources.unix. Copyright info for that program
 * is included below as required.
 * -----
 * - Copyright (c) Ian F. Darwin, 1987. Written by Ian F. Darwin.
 *
 * This software is not subject to any license of the American Telephone and
 * Telegraph Company or of the Regents of the University of California.
 *
 * Permission is granted to anyone to use this software for any purpose on any
 * computer system, and to alter it and redistribute it freely, subject to
 * the following restrictions:
 *
 * 1. The author is not responsible for the consequences of use of this
 * software, no matter how awful, even if they arise from flaws in it.
 *
 * 2. The origin of this software must not be misrepresented, either by
 * explicit claim or by omission. Since few users ever read sources, credits
 * must appear in the documentation.
 *
 * 3. Altered versions must be plainly marked as such, and must not be
 * misrepresented as being the original software. Since few users ever read
 * sources, credits must appear in the documentation.
 *
 * 4. This notice may not be removed or altered.
```

### modules\mappers\mod\_imap.c component

"macmartinized" polygon code copyright 1992 by Eric Haines, erich@eye.com

### server\util\_md5.c

NCSA HTTPd Server  
Software Development Group  
National Center for Supercomputing Applications  
University of Illinois at Urbana-Champaign  
605 E. Springfield, Champaign, IL 61820

httpd@ncsa.uiuc.edu

Copyright (C) 1995, Board of Trustees of the University of Illinois

md5.c: NCSA HTTPd code which uses the md5c.c RSA Code

Original Code Copyright (C) 1994, Jeff Hostetler, Spyglass, Inc.  
Portions of Content-MD5 code Copyright (C) 1993, 1994 by Carnegie Mellon University (see Copyright below).

Portions of Content-MD5 code Copyright (C) 1991 Bell Communications Research, Inc. (Bellcore) (see Copyright below).

Portions extracted from mpack, John G. Myers - jgm+@cmu.edu

Content-MD5 Code contributed by Martin Hamilton (martin@net.lut.ac.uk)

/ these portions extracted from mpack, John G. Myers - jgm+@cmu.edu /  
/ (C) Copyright 1993,1994 by Carnegie Mellon University  
All Rights Reserved.

Permission to use, copy, modify, distribute, and sell this software and its documentation for any purpose is hereby granted without fee, provided that the above copyright notice appear in all copies and that both that copyright notice and this permission notice appear in supporting documentation, and that the name of Carnegie Mellon University not be used in advertising or publicity pertaining to distribution of the software without specific, written prior permission. Carnegie Mellon University makes no representations about the suitability of this software for any purpose. It is provided "as is" without express or implied warranty.

CARNEGIE MELLON UNIVERSITY DISCLAIMS ALL WARRANTIES WITH REGARD TO THIS SOFTWARE, INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS, IN NO EVENT SHALL CARNEGIE MELLON UNIVERSITY BE LIABLE FOR ANY SPECIAL, INDIRECT OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

Copyright (c) 1991 Bell Communications Research, Inc. (Bellcore)

Permission to use, copy, modify, and distribute this material for any purpose and without fee is hereby granted, provided that the above copyright notice and this permission notice appear in all copies, and that the name of Bellcore not be used in advertising or publicity pertaining to this material without the specific, prior written permission of an authorized representative of Bellcore. BELLCORE MAKES NO REPRESENTATIONS ABOUT THE ACCURACY OR SUITABILITY OF THIS MATERIAL FOR ANY PURPOSE. IT IS PROVIDED "AS IS", WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES.

## src/lib\apr\include\apr\_md5.h component

This work is derived from material Copyright RSA Data Security, Inc.

The RSA copyright statement and Licence for that original material is included below. This is followed by the Apache copyright statement and licence for the modifications made to that material.

/ Copyright (C) 1991-2, RSA Data Security, Inc. Created 1991. All rights reserved.

License to copy and use this software is granted provided that it is identified as the "RSA Data Security, Inc. MD5 Message-Digest

Algorithm" in all material mentioning or referencing this software or this function.

License is also granted to make and use derivative works provided that such works are identified as "derived from the RSA Data Security, Inc. MD5 Message-Digest Algorithm" in all material mentioning or referencing the derived work.

RSA Data Security, Inc. makes no representations concerning either the merchantability of this software or the suitability of this software for any particular purpose. It is provided "as is" without express or implied warranty of any kind.

These notices must be retained in any copies of any part of this documentation and/or software.

## srclib\apr\passwd\apr\_md5.c Component

For the srclib\apr\passwd\apr\_md5.c component:

/ This is work is derived from material Copyright RSA Data Security, Inc.

The RSA copyright statement and Licence for that original material is included below. This is followed by the Apache copyright statement and licence for the modifications made to that material.

/ MD5C.C - RSA Data Security, Inc., MD5 message-digest algorithm  
/

/ Copyright (C) 1991-2, RSA Data Security, Inc. Created 1991. All rights reserved.

License to copy and use this software is granted provided that it is identified as the "RSA Data Security, Inc. MD5 Message-Digest Algorithm" in all material mentioning or referencing this software or this function.

License is also granted to make and use derivative works provided that such works are identified as "derived from the RSA Data Security, Inc. MD5 Message-Digest Algorithm" in all material mentioning or referencing the derived work.

RSA Data Security, Inc. makes no representations concerning either the merchantability of this software or the suitability of this software for any particular purpose. It is provided "as is" without express or implied warranty of any kind.

These notices must be retained in any copies of any part of this documentation and/or software.

## apr\_md5\_encode() Routine Component

The apr\_md5\_encode() routine uses much code obtained from the FreeBSD 3.0 MD5 crypt() function, which is licenced as follows:

-----  
"THE BEER-WARE LICENSE" (Revision 42):  
<phk@login.dknet.dk wrote this file. As long as you retain this notice you can do whatever you want with this stuff. If we meet some day, and you think this stuff is worth it, you can buy me a beer in return. Poul-Henning Kamp  
-----

## srclib\apr-util\crypto\apr\_md4.c Component

This is derived from material copyright RSA Data Security, Inc. Their notice is reproduced below in its entirety.

Copyright (C) 1991-2, RSA Data Security, Inc. Created 1991. All rights reserved.

License to copy and use this software is granted provided that it is identified as the "RSA Data Security, Inc. MD4 Message-Digest Algorithm" in all material mentioning or referencing this software or this function.

License is also granted to make and use derivative works provided that such works are identified as "derived from the RSA Data Security, Inc. MD4 Message-Digest Algorithm" in all material mentioning or referencing the derived work.

RSA Data Security, Inc. makes no representations concerning either the merchantability of this software or the suitability of this software for any particular purpose. It is provided "as is" without express or implied warranty of any kind.

These notices must be retained in any copies of any part of this documentation and/or software.

### src\lib\apr-util\include\apr\_md4.h Component

This is derived from material copyright RSA Data Security, Inc. Their notice is reproduced below in its entirety.

Copyright (C) 1991-2, RSA Data Security, Inc. Created 1991. All rights reserved.

License to copy and use this software is granted provided that it is identified as the "RSA Data Security, Inc. MD4 Message-Digest Algorithm" in all material mentioning or referencing this software or this function.

License is also granted to make and use derivative works provided that such works are identified as "derived from the RSA Data Security, Inc. MD4 Message-Digest Algorithm" in all material mentioning or referencing the derived work.

RSA Data Security, Inc. makes no representations concerning either the merchantability of this software or the suitability of this software for any particular purpose. It is provided "as is" without express or implied warranty of any kind.

These notices must be retained in any copies of any part of this documentation and/or software.

### src\lib\apr-util\test\testdbm.c Component

The Apache Software License, Version 1.1

Copyright (c) 2000-2002 The Apache Software Foundation. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the



distribution.

3. The end-user documentation included with the redistribution, if any, must include the following acknowledgment:  
"This product includes software developed by the Apache Software Foundation (<http://www.apache.org/>)."  
Alternately, this acknowledgment may appear in the software itself, if and wherever such third-party acknowledgments normally appear.

4. The names "Apache" and "Apache Software Foundation" must not be used to endorse or promote products derived from this software without prior written permission. For written permission, please contact [apache@apache.org](mailto:apache@apache.org).

5. Products derived from this software may not be called "Apache", nor may "Apache" appear in their name, without prior written permission of the Apache Software Foundation.

THIS SOFTWARE IS PROVIDED ``AS IS'' AND ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE APACHE SOFTWARE FOUNDATION OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

=====

This software consists of voluntary contributions made by many individuals on behalf of the Apache Software Foundation. For more information on the Apache Software Foundation, please see <http://www.apache.org/>.

This file came from the SDBM package (written by [oz@nexus.yorku.ca](mailto:oz@nexus.yorku.ca)). That package was under public domain. This file has been ported to APR, updated to ANSI C and other, newer idioms, and added to the Apache codebase under the above copyright and license.

## src/lib\apr-util\test\testmd4.c Component

This is derived from material copyright RSA Data Security, Inc. Their notice is reproduced below in its entirety.

Copyright (C) 1990-2, RSA Data Security, Inc. Created 1990. All rights reserved.

RSA Data Security, Inc. makes no representations concerning either the merchantability of this software or the suitability of this software for any particular purpose. It is provided "as is" without express or implied warranty of any kind.

These notices must be retained in any copies of any part of this documentation and/or software.

## src/lib\apr-util\xml\expat\conf\tools\install-sh Component

install - install a program, script, or datafile  
This comes from X11R5 (mit/util/scripts/install.sh).

Copyright 1991 by the Massachusetts Institute of Technology

Permission to use, copy, modify, distribute, and sell this software and its

documentation for any purpose is hereby granted without fee, provided that the above copyright notice appear in all copies and that both that copyright notice and this permission notice appear in supporting documentation, and that the name of M.I.T. not be used in advertising or publicity pertaining to distribution of the software without specific, written prior permission. M.I.T. makes no representations about the suitability of this software for any purpose. It is provided "as is" without express or implied warranty.

## PCRE Component

PCRE is a library of functions to support regular expressions whose syntax and semantics are as close as possible to those of the Perl 5 language.

Written by: Philip Hazel ph10@cam.ac.uk

University of Cambridge Computing Service,  
Cambridge, England. Phone: +44 1223 334714.

Copyright (c) 1997-2001 University of Cambridge

Permission is granted to anyone to use this software for any purpose on any computer system, and to redistribute it freely, subject to the following restrictions:

1. This software is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.
2. The origin of this software must not be misrepresented, either by explicit claim or by omission. In practice, this means that if you use PCRE in software which you distribute to others, commercially or otherwise, you must put a sentence like this

Regular expression support is provided by the PCRE library package, which is open source software, written by Philip Hazel, and copyright by the University of Cambridge, England.

somewhere reasonably visible in your documentation and in any relevant files or online help data or similar. A reference to the ftp site for the source, that is, to

<ftp://ftp.csx.cam.ac.uk/pub/software/programming/pcre/>

should also be given in the documentation.

3. Altered versions must be plainly marked as such, and must not be misrepresented as being the original software.
4. If PCRE is embedded in any software that is released under the GNU General Purpose Licence (GPL), or Lesser General Purpose Licence (LGPL), then the terms of that licence shall supersede any condition above with which it is incompatible.

The documentation for PCRE, supplied in the "doc" directory, is distributed under the same terms as the software itself.

## test\zb.c Component

ZeusBench V1.01  
=====

This program is Copyright (C) Zeus Technology Limited 1996.

This program may be used and copied freely providing this copyright notice is not removed.

This software is provided "as is" and any express or implied warranties, including but not limited to, the implied warranties of merchantability and fitness for a particular purpose are disclaimed. In no event shall Zeus Technology Ltd. be liable for any direct, indirect, incidental, special, exemplary, or consequential damages (including, but not limited to, procurement of substitute good or services; loss of use, data, or profits; or business interruption) however caused and on theory of liability. Whether in contract, strict liability or tort (including negligence or otherwise) arising in any way out of the use of this software, even if advised of the possibility of such damage.

Written by Adam Twiss (adam@zeus.co.uk). March 1996

Thanks to the following people for their input:  
Mike Belshe (mbelshe@netscape.com)  
Michael Campanella (campanella@stevms.enet.dec.com)

## expat xml parser Component

Copyright (c) 1998, 1999, 2000 Thai Open Source Software Center Ltd  
and Clark Cooper

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

## Cyrus IMAP License

\* Copyright (c) 1994-2000 Carnegie Mellon University. All rights reserved.  
\*  
\* Redistribution and use in source and binary forms, with or without  
\* modification, are permitted provided that the following conditions  
\* are met:  
\*  
\* 1. Redistributions of source code must retain the above copyright  
\* notice, this list of conditions and the following disclaimer.  
\*  
\* 2. Redistributions in binary form must reproduce the above copyright  
\* notice, this list of conditions and the following disclaimer in  
\* the documentation and/or other materials provided with the  
\* distribution.  
\*  
\* 3. The name "Carnegie Mellon University" must not be used to  
\* endorse or promote products derived from this software without  
\* prior written permission. For permission or any legal  
\* details, please contact  
\* Office of Technology Transfer  
\* Carnegie Mellon University  
\* 5000 Forbes Avenue  
\* Pittsburgh, PA 15213-3890  
\* (412) 268-4387, fax: (412) 268-7395  
\* tech-transfer@andrew.cmu.edu

\*  
 \* 4. Redistributions of any form whatsoever must retain the following  
 \* acknowledgment:  
 \* "This product includes software developed by Computing Services  
 \* at Carnegie Mellon University (<http://www.cmu.edu/computing/>)."  
 \*  
 \* CARNEGIE MELLON UNIVERSITY DISCLAIMS ALL WARRANTIES WITH REGARD TO  
 \* THIS SOFTWARE, INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY  
 \* AND FITNESS, IN NO EVENT SHALL CARNEGIE MELLON UNIVERSITY BE LIABLE  
 \* FOR ANY SPECIAL, INDIRECT OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES  
 \* WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN  
 \* AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING  
 \* OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

If you find this software useful and valuable in your work, we would welcome any support you can offer toward continuing this work.

We gratefully accept contributions, whether intellectual or monetary. Intellectual contributions in the form of code or constructive collaboration can be directed to [cyrus-bugs+@andrew.cmu.edu](mailto:cyrus-bugs+@andrew.cmu.edu) (even if it is not a bug).

If you wish to provide financial support to the Cyrus Project, send a check payable to "Carnegie Mellon University" to

Project Cyrus  
 Computing Services  
 Carnegie Mellon University  
 5000 Forbes Ave  
 Pittsburgh, PA 15213  
 USA

## Mozilla License

### MOZILLA PUBLIC LICENSE Version 1.1

-----

#### 1. Definitions.

1.0.1. "Commercial Use" means distribution or otherwise making the Covered Code available to a third party.

1.1. "Contributor" means each entity that creates or contributes to the creation of Modifications.

1.2. "Contributor Version" means the combination of the Original Code, prior Modifications used by a Contributor, and the Modifications made by that particular Contributor.

1.3. "Covered Code" means the Original Code or Modifications or the combination of the Original Code and Modifications, in each case including portions thereof.

1.4. "Electronic Distribution Mechanism" means a mechanism generally accepted in the software development community for the electronic transfer of data.

1.5. "Executable" means Covered Code in any form other than Source Code.

1.6. "Initial Developer" means the individual or entity identified as the Initial Developer in the Source Code notice required by Exhibit A.

1.7. "Larger Work" means a work which combines Covered Code or portions thereof with code not governed by the terms of this License.

1.8. "License" means this document.

1.8.1. "Licensable" means having the right to grant, to the maximum extent possible, whether at the time of the initial grant or subsequently acquired, any and all of the rights conveyed herein.

1.9. "Modifications" means any addition to or deletion from the substance or structure of either the Original Code or any previous Modifications. When Covered Code is released as a series of files, a Modification is:

A. Any addition to or deletion from the contents of a file containing Original Code or previous Modifications.

B. Any new file that contains any part of the Original Code or previous Modifications.

1.10. "Original Code" means Source Code of computer software code which is described in the Source Code notice required by Exhibit A as Original Code, and which, at the time of its release under this License is not already Covered Code governed by this License.

1.10.1. "Patent Claims" means any patent claim(s), now owned or hereafter acquired, including without limitation, method, process, and apparatus claims, in any patent Licensable by grantor.

1.11. "Source Code" means the preferred form of the Covered Code for making modifications to it, including all modules it contains, plus any associated interface definition files, scripts used to control compilation and installation of an Executable, or source code differential comparisons against either the Original Code or another well known, available Covered Code of the Contributor's choice. The Source Code can be in a compressed or archival form, provided the appropriate decompression or de-archiving software is widely available for no charge.

1.12. "You" (or "Your") means an individual or a legal entity exercising rights under, and complying with all of the terms of, this License or a future version of this License issued under Section 6.1. For legal entities, "You" includes any entity which controls, is controlled by, or is under common control with You. For purposes of this definition, "control" means (a) the power, direct or indirect, to cause the direction or management of such entity, whether by contract or otherwise, or (b) ownership of more than fifty percent (50%) of the outstanding shares or beneficial ownership of such entity.

## 2. Source Code License.

### 2.1. The Initial Developer Grant.

The Initial Developer hereby grants You a world-wide, royalty-free, non-exclusive license, subject to third party intellectual property claims:

(a) under intellectual property rights (other than patent or trademark) Licensable by Initial Developer to use, reproduce, modify, display, perform, sublicense and distribute the Original Code (or portions thereof) with or without Modifications, and/or as part of a Larger Work; and

(b) under Patents Claims infringed by the making, using or selling of Original Code, to make, have made, use, practice, sell, and offer for sale, and/or otherwise dispose of the

Original Code (or portions thereof).

(c) the licenses granted in this Section 2.1(a) and (b) are effective on the date Initial Developer first distributes Original Code under the terms of this License.

(d) Notwithstanding Section 2.1(b) above, no patent license is granted: 1) for code that You delete from the Original Code; 2) separate from the Original Code; or 3) for infringements caused by: i) the modification of the Original Code or ii) the combination of the Original Code with other software or devices.

## 2.2. Contributor Grant.

Subject to third party intellectual property claims, each Contributor hereby grants You a world-wide, royalty-free, non-exclusive license

(a) under intellectual property rights (other than patent or trademark) Licensable by Contributor, to use, reproduce, modify, display, perform, sublicense and distribute the Modifications created by such Contributor (or portions thereof) either on an unmodified basis, with other Modifications, as Covered Code and/or as part of a Larger Work; and

(b) under Patent Claims infringed by the making, using, or selling of Modifications made by that Contributor either alone and/or in combination with its Contributor Version (or portions of such combination), to make, use, sell, offer for sale, have made, and/or otherwise dispose of: 1) Modifications made by that Contributor (or portions thereof); and 2) the combination of Modifications made by that Contributor with its Contributor Version (or portions of such combination).

(c) the licenses granted in Sections 2.2(a) and 2.2(b) are effective on the date Contributor first makes Commercial Use of the Covered Code.

(d) Notwithstanding Section 2.2(b) above, no patent license is granted: 1) for any code that Contributor has deleted from the Contributor Version; 2) separate from the Contributor Version; 3) for infringements caused by: i) third party modifications of Contributor Version or ii) the combination of Modifications made by that Contributor with other software (except as part of the Contributor Version) or other devices; or 4) under Patent Claims infringed by Covered Code in the absence of Modifications made by that Contributor.

## 3. Distribution Obligations.

### 3.1. Application of License.

The Modifications which You create or to which You contribute are governed by the terms of this License, including without limitation Section 2.2. The Source Code version of Covered Code may be distributed only under the terms of this License or a future version of this License released under Section 6.1, and You must include a copy of this License with every copy of the Source Code You distribute. You may not offer or impose any terms on any Source Code version that alters or restricts the applicable version of this License or the recipients' rights hereunder. However, You may include an additional document offering the additional rights described in Section 3.5.

### 3.2. Availability of Source Code.

Any Modification which You create or to which You contribute must be made available in Source Code form under the terms of this License either on the same media as an Executable version or via an accepted

Electronic Distribution Mechanism to anyone to whom you made an Executable version available; and if made available via Electronic Distribution Mechanism, must remain available for at least twelve (12) months after the date it initially became available, or at least six (6) months after a subsequent version of that particular Modification has been made available to such recipients. You are responsible for ensuring that the Source Code version remains available even if the Electronic Distribution Mechanism is maintained by a third party.

### 3.3. Description of Modifications.

You must cause all Covered Code to which You contribute to contain a file documenting the changes You made to create that Covered Code and the date of any change. You must include a prominent statement that the Modification is derived, directly or indirectly, from Original Code provided by the Initial Developer and including the name of the Initial Developer in (a) the Source Code, and (b) in any notice in an Executable version or related documentation in which You describe the origin or ownership of the Covered Code.

### 3.4. Intellectual Property Matters

#### (a) Third Party Claims.

If Contributor has knowledge that a license under a third party's intellectual property rights is required to exercise the rights granted by such Contributor under Sections 2.1 or 2.2, Contributor must include a text file with the Source Code distribution titled "LEGAL" which describes the claim and the party making the claim in sufficient detail that a recipient will know whom to contact. If Contributor obtains such knowledge after the Modification is made available as described in Section 3.2, Contributor shall promptly modify the LEGAL file in all copies Contributor makes available thereafter and shall take other steps (such as notifying appropriate mailing lists or newsgroups) reasonably calculated to inform those who received the Covered Code that new knowledge has been obtained.

#### (b) Contributor APIs.

If Contributor's Modifications include an application programming interface and Contributor has knowledge of patent licenses which are reasonably necessary to implement that API, Contributor must also include this information in the LEGAL file.

#### (c) Representations.

Contributor represents that, except as disclosed pursuant to Section 3.4(a) above, Contributor believes that Contributor's Modifications are Contributor's original creation(s) and/or Contributor has sufficient rights to grant the rights conveyed by this License.

### 3.5. Required Notices.

You must duplicate the notice in Exhibit A in each file of the Source Code. If it is not possible to put such notice in a particular Source Code file due to its structure, then You must include such notice in a location (such as a relevant directory) where a user would be likely to look for such a notice. If You created one or more Modification(s) You may add your name as a Contributor to the notice described in Exhibit A. You must also duplicate this License in any documentation for the Source Code where You describe recipients' rights or ownership rights relating to Covered Code. You may choose to offer, and to charge a fee for, warranty, support, indemnity or liability obligations to one or more recipients of Covered Code. However, You may do so only on Your own behalf, and not on behalf of the Initial Developer or any Contributor. You must make it absolutely clear than any such warranty, support, indemnity or liability obligation is offered by You alone, and You hereby agree to indemnify the Initial Developer and every Contributor for any liability incurred by the

Initial Developer or such Contributor as a result of warranty, support, indemnity or liability terms You offer.

### 3.6. Distribution of Executable Versions.

You may distribute Covered Code in Executable form only if the requirements of Section 3.1-3.5 have been met for that Covered Code, and if You include a notice stating that the Source Code version of the Covered Code is available under the terms of this License, including a description of how and where You have fulfilled the obligations of Section 3.2. The notice must be conspicuously included in any notice in an Executable version, related documentation or collateral in which You describe recipients' rights relating to the Covered Code. You may distribute the Executable version of Covered Code or ownership rights under a license of Your choice, which may contain terms different from this License, provided that You are in compliance with the terms of this License and that the license for the Executable version does not attempt to limit or alter the recipient's rights in the Source Code version from the rights set forth in this License. If You distribute the Executable version under a different license You must make it absolutely clear that any terms which differ from this License are offered by You alone, not by the Initial Developer or any Contributor. You hereby agree to indemnify the Initial Developer and every Contributor for any liability incurred by the Initial Developer or such Contributor as a result of any such terms You offer.

### 3.7. Larger Works.

You may create a Larger Work by combining Covered Code with other code not governed by the terms of this License and distribute the Larger Work as a single product. In such a case, You must make sure the requirements of this License are fulfilled for the Covered Code.

## 4. Inability to Comply Due to Statute or Regulation.

If it is impossible for You to comply with any of the terms of this License with respect to some or all of the Covered Code due to statute, judicial order, or regulation then You must: (a) comply with the terms of this License to the maximum extent possible; and (b) describe the limitations and the code they affect. Such description must be included in the LEGAL file described in Section 3.4 and must be included with all distributions of the Source Code. Except to the extent prohibited by statute or regulation, such description must be sufficiently detailed for a recipient of ordinary skill to be able to understand it.

## 5. Application of this License.

This License applies to code to which the Initial Developer has attached the notice in Exhibit A and to related Covered Code.

## 6. Versions of the License.

### 6.1. New Versions.

Netscape Communications Corporation ("Netscape") may publish revised and/or new versions of the License from time to time. Each version will be given a distinguishing version number.

### 6.2. Effect of New Versions.

Once Covered Code has been published under a particular version of the License, You may always continue to use it under the terms of that version. You may also choose to use such Covered Code under the terms of any subsequent version of the License published by Netscape. No one other than Netscape has the right to modify the terms applicable to Covered Code created under this License.



### 6.3. Derivative Works.

If You create or use a modified version of this License (which you may only do in order to apply it to code which is not already Covered Code governed by this License), You must (a) rename Your license so that the phrases "Mozilla", "MOZILLAPL", "MOZPL", "Netscape", "MPL", "NPL" or any confusingly similar phrase do not appear in your license (except to note that your license differs from this License) and (b) otherwise make it clear that Your version of the license contains terms which differ from the Mozilla Public License and Netscape Public License. (Filling in the name of the Initial Developer, Original Code or Contributor in the notice described in Exhibit A shall not of themselves be deemed to be modifications of this License.)

## 7. DISCLAIMER OF WARRANTY.

COVERED CODE IS PROVIDED UNDER THIS LICENSE ON AN "AS IS" BASIS, WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, WARRANTIES THAT THE COVERED CODE IS FREE OF DEFECTS, MERCHANTABILITY, FIT FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE COVERED CODE IS WITH YOU. SHOULD ANY COVERED CODE PROVE DEFECTIVE IN ANY RESPECT, YOU (NOT THE INITIAL DEVELOPER OR ANY OTHER CONTRIBUTOR) ASSUME THE COST OF ANY NECESSARY SERVICING, REPAIR OR CORRECTION. THIS DISCLAIMER OF WARRANTY CONSTITUTES AN ESSENTIAL PART OF THIS LICENSE. NO USE OF ANY COVERED CODE IS AUTHORIZED HEREUNDER EXCEPT UNDER THIS DISCLAIMER.

## 8. TERMINATION.

8.1. This License and the rights granted hereunder will terminate automatically if You fail to comply with terms herein and fail to cure such breach within 30 days of becoming aware of the breach. All sublicenses to the Covered Code which are properly granted shall survive any termination of this License. Provisions which, by their nature, must remain in effect beyond the termination of this License shall survive.

8.2. If You initiate litigation by asserting a patent infringement claim (excluding declaratory judgment actions) against Initial Developer or a Contributor (the Initial Developer or Contributor against whom You file such action is referred to as "Participant") alleging that:

(a) such Participant's Contributor Version directly or indirectly infringes any patent, then any and all rights granted by such Participant to You under Sections 2.1 and/or 2.2 of this License shall, upon 60 days notice from Participant terminate prospectively, unless if within 60 days after receipt of notice You either: (i) agree in writing to pay Participant a mutually agreeable reasonable royalty for Your past and future use of Modifications made by such Participant, or (ii) withdraw Your litigation claim with respect to the Contributor Version against such Participant. If within 60 days of notice, a reasonable royalty and payment arrangement are not mutually agreed upon in writing by the parties or the litigation claim is not withdrawn, the rights granted by Participant to You under Sections 2.1 and/or 2.2 automatically terminate at the expiration of the 60 day notice period specified above.

(b) any software, hardware, or device, other than such Participant's Contributor Version, directly or indirectly infringes any patent, then any rights granted to You by such Participant under Sections 2.1(b) and 2.2(b) are revoked effective as of the date You first made, used, sold, distributed, or had made, Modifications made by that Participant.

8.3. If You assert a patent infringement claim against Participant

alleging that such Participant's Contributor Version directly or indirectly infringes any patent where such claim is resolved (such as by license or settlement) prior to the initiation of patent infringement litigation, then the reasonable value of the licenses granted by such Participant under Sections 2.1 or 2.2 shall be taken into account in determining the amount or value of any payment or license.

8.4. In the event of termination under Sections 8.1 or 8.2 above, all end user license agreements (excluding distributors and resellers) which have been validly granted by You or any distributor hereunder prior to termination shall survive termination.

#### 9. LIMITATION OF LIABILITY.

UNDER NO CIRCUMSTANCES AND UNDER NO LEGAL THEORY, WHETHER TORT (INCLUDING NEGLIGENCE), CONTRACT, OR OTHERWISE, SHALL YOU, THE INITIAL DEVELOPER, ANY OTHER CONTRIBUTOR, OR ANY DISTRIBUTOR OF COVERED CODE, OR ANY SUPPLIER OF ANY OF SUCH PARTIES, BE LIABLE TO ANY PERSON FOR ANY INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES OF ANY CHARACTER INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF GOODWILL, WORK STOPPAGE, COMPUTER FAILURE OR MALFUNCTION, OR ANY AND ALL OTHER COMMERCIAL DAMAGES OR LOSSES, EVEN IF SUCH PARTY SHALL HAVE BEEN INFORMED OF THE POSSIBILITY OF SUCH DAMAGES. THIS LIMITATION OF LIABILITY SHALL NOT APPLY TO LIABILITY FOR DEATH OR PERSONAL INJURY RESULTING FROM SUCH PARTY'S NEGLIGENCE TO THE EXTENT APPLICABLE LAW PROHIBITS SUCH LIMITATION. SOME JURISDICTIONS DO NOT ALLOW THE EXCLUSION OR LIMITATION OF INCIDENTAL OR CONSEQUENTIAL DAMAGES, SO THIS EXCLUSION AND LIMITATION MAY NOT APPLY TO YOU.

#### 10. U.S. GOVERNMENT END USERS.

The Covered Code is a "commercial item," as that term is defined in 48 C.F.R. 2.101 (Oct. 1995), consisting of "commercial computer software" and "commercial computer software documentation," as such terms are used in 48 C.F.R. 12.212 (Sept. 1995). Consistent with 48 C.F.R. 12.212 and 48 C.F.R. 227.7202-1 through 227.7202-4 (June 1995), all U.S. Government End Users acquire Covered Code with only those rights set forth herein.

#### 11. MISCELLANEOUS.

This License represents the complete agreement concerning subject matter hereof. If any provision of this License is held to be unenforceable, such provision shall be reformed only to the extent necessary to make it enforceable. This License shall be governed by California law provisions (except to the extent applicable law, if any, provides otherwise), excluding its conflict-of-law provisions. With respect to disputes in which at least one party is a citizen of, or an entity chartered or registered to do business in the United States of America, any litigation relating to this License shall be subject to the jurisdiction of the Federal Courts of the Northern District of California, with venue lying in Santa Clara County, California, with the losing party responsible for costs, including without limitation, court costs and reasonable attorneys' fees and expenses. The application of the United Nations Convention on Contracts for the International Sale of Goods is expressly excluded. Any law or regulation which provides that the language of a contract shall be construed against the drafter shall not apply to this License.

#### 12. RESPONSIBILITY FOR CLAIMS.

As between Initial Developer and the Contributors, each party is responsible for claims and damages arising, directly or indirectly,

out of its utilization of rights under this License and You agree to work with Initial Developer and Contributors to distribute such responsibility on an equitable basis. Nothing herein is intended or shall be deemed to constitute any admission of liability.

### 13. MULTIPLE-LICENSED CODE.

Initial Developer may designate portions of the Covered Code as "Multiple-Licensed". "Multiple-Licensed" means that the Initial Developer permits you to utilize portions of the Covered Code under Your choice of the NPL or the alternative licenses, if any, specified by the Initial Developer in the file described in Exhibit A.

#### EXHIBIT A -Mozilla Public License.

``The contents of this file are subject to the Mozilla Public License Version 1.1 (the "License"); you may not use this file except in compliance with the License. You may obtain a copy of the License at <http://www.mozilla.org/MPL/>

Software distributed under the License is distributed on an "AS IS" basis, WITHOUT WARRANTY OF ANY KIND, either express or implied. See the License for the specific language governing rights and limitations under the License.

The Original Code is \_\_\_\_\_.

The Initial Developer of the Original Code is \_\_\_\_\_.  
Portions created by \_\_\_\_\_ are Copyright (C) \_\_\_\_\_  
\_\_\_\_\_. All Rights Reserved.

Contributor(s): \_\_\_\_\_.

Alternatively, the contents of this file may be used under the terms of the \_\_\_\_\_ license (the "[\_\_\_] License"), in which case the provisions of [\_\_\_\_\_] License are applicable instead of those above. If you wish to allow use of your version of this file only under the terms of the [\_\_\_] License and not to allow others to use your version of this file under the MPL, indicate your decision by deleting the provisions above and replace them with the notice and other provisions required by the [\_\_\_] License. If you do not delete the provisions above, a recipient may use your version of this file under either the MPL or the [\_\_\_] License."

[NOTE: The text of this Exhibit A may differ slightly from the text of the notices in the Source Code files of the Original Code. You should use the text of this Exhibit A rather than the text found in the Original Code Source Code for Your Modifications.]

## OpenLDAP License

Copyright 1998-2004 The OpenLDAP Foundation  
All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted only as authorized by the OpenLDAP Public License.

A copy of this license is available in the file LICENSE in the top-level directory of the distribution or, alternatively, at

< <http://www.OpenLDAP.org/license.html>.

OpenLDAP is a registered trademark of the OpenLDAP Foundation.

Individual files and/or contributed packages may be copyright by other parties and subject to additional restrictions.

This work is derived from the University of Michigan LDAP v3.3 distribution. Information concerning this software is available at < <http://www.umich.edu/~dirsvcs/ldap/>.

This work also contains materials derived from public sources.

Additional information about OpenLDAP can be obtained at < <http://www.openldap.org/>.

---

Portions Copyright 1998-2004 Kurt D. Zeilenga.  
Portions Copyright 1998-2004 Net Boolean Incorporated.  
Portions Copyright 2001-2004 IBM Corporation.  
All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted only as authorized by the OpenLDAP Public License.

---

Portions Copyright 1999-2003 Howard Y.H. Chu.  
Portions Copyright 1999-2003 Symas Corporation.  
Portions Copyright 1998-2003 Hallvard B. Furuseth.  
All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that this notice is preserved. The names of the copyright holders may not be used to endorse or promote products derived from this software without their specific prior written permission. This software is provided ``as is'' without express or implied warranty.

---

Portions Copyright (c) 1992-1996 Regents of the University of Michigan.  
All rights reserved.

Redistribution and use in source and binary forms are permitted provided that this notice is preserved and that due credit is given to the University of Michigan at Ann Arbor. The name of the University may not be used to endorse or promote products derived from this software without specific prior written permission. This software is provided ``as is'' without express or implied warranty.

## Sendmail License

### SENDMAIL LICENSE

The following license terms and conditions apply, unless a different license is obtained from Sendmail, Inc., 6425 Christie Ave, Fourth Floor, Emeryville, CA 94608, USA or by electronic mail at [license@sendmail.com](mailto:license@sendmail.com).

License Terms:

Use, Modification and Redistribution (including distribution of any modified or derived work) in source and binary forms is permitted only if each of the following conditions is met:

1. Redistributions qualify as "freeware" or "Open Source Software" under one of the following terms:
  - (a) Redistributions are made at no charge beyond the reasonable cost of

materials and delivery.

- (b) Redistributions are accompanied by a copy of the Source Code or by an irrevocable offer to provide a copy of the Source Code for up to three years at the cost of materials and delivery. Such redistributions must allow further use, modification, and redistribution of the Source Code under substantially the same terms as this license. For the purposes of redistribution "Source Code" means the complete compilable and linkable source code of sendmail including all modifications.
2. Redistributions of source code must retain the copyright notices as they appear in each source code file, these license terms, and the disclaimer/limitation of liability set forth as paragraph 6 below.
  3. Redistributions in binary form must reproduce the Copyright Notice, these license terms, and the disclaimer/limitation of liability set forth as paragraph 6 below, in the documentation and/or other materials provided with the distribution. For the purposes of binary distribution the "Copyright Notice" refers to the following language:  
"Copyright (c) 1998-2004 Sendmail, Inc. All rights reserved."
  4. Neither the name of Sendmail, Inc. nor the University of California nor the names of their contributors may be used to endorse or promote products derived from this software without specific prior written permission. The name "sendmail" is a trademark of Sendmail, Inc.
  5. All redistributions must comply with the conditions imposed by the University of California on certain embedded code, whose copyright notice and conditions for redistribution are as follows:
    - (a) Copyright (c) 1988, 1993 The Regents of the University of California. All rights reserved.
    - (b) Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:
      - (i) Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
      - (ii) Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
      - (iii) Neither the name of the University nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.
  6. Disclaimer/Limitation of Liability: THIS SOFTWARE IS PROVIDED BY SENDMAIL, INC. AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL SENDMAIL, INC., THE REGENTS OF THE UNIVERSITY OF CALIFORNIA OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

\$Revision: 8.13 \$, Last updated \$Date: 2004/05/11 23:57:57 \$

## Tomcat License

Apache License  
Version 2.0, January 2004  
<http://www.apache.org/licenses/>

TERMS AND CONDITIONS FOR USE, REPRODUCTION, AND DISTRIBUTION

1. Definitions.

"License" shall mean the terms and conditions for use, reproduction, and distribution as defined by Sections 1 through 9 of this document.

"Licensor" shall mean the copyright owner or entity authorized by the copyright owner that is granting the License.

"Legal Entity" shall mean the union of the acting entity and all other entities that control, are controlled by, or are under common control with that entity. For the purposes of this definition, "control" means (i) the power, direct or indirect, to cause the direction or management of such entity, whether by contract or otherwise, or (ii) ownership of fifty percent (50%) or more of the outstanding shares, or (iii) beneficial ownership of such entity.

"You" (or "Your") shall mean an individual or Legal Entity exercising permissions granted by this License.

"Source" form shall mean the preferred form for making modifications, including but not limited to software source code, documentation source, and configuration files.

"Object" form shall mean any form resulting from mechanical transformation or translation of a Source form, including but not limited to compiled object code, generated documentation, and conversions to other media types.

"Work" shall mean the work of authorship, whether in Source or Object form, made available under the License, as indicated by a copyright notice that is included in or attached to the work (an example is provided in the Appendix below).

"Derivative Works" shall mean any work, whether in Source or Object form, that is based on (or derived from) the Work and for which the editorial revisions, annotations, elaborations, or other modifications represent, as a whole, an original work of authorship. For the purposes of this License, Derivative Works shall not include works that remain separable from, or merely link (or bind by name) to the interfaces of, the Work and Derivative Works thereof.

"Contribution" shall mean any work of authorship, including the original version of the Work and any modifications or additions to that Work or Derivative Works thereof, that is intentionally submitted to Licensor for inclusion in the Work by the copyright owner or by an individual or Legal Entity authorized to submit on behalf of the copyright owner. For the purposes of this definition, "submitted" means any form of electronic, verbal, or written communication sent to the Licensor or its representatives, including but not limited to communication on electronic mailing lists, source code control systems, and issue tracking systems that are managed by, or on behalf of, the Licensor for the purpose of discussing and improving the Work, but excluding communication that is conspicuously marked or otherwise designated in writing by the copyright owner as "Not a Contribution."

"Contributor" shall mean Licensor and any individual or Legal Entity on behalf of whom a Contribution has been received by Licensor and subsequently incorporated within the Work.

2. Grant of Copyright License. Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable copyright license to reproduce, prepare Derivative Works of, publicly display, publicly perform, sublicense, and distribute the Work and such Derivative Works in Source or Object form.
3. Grant of Patent License. Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable (except as stated in this section) patent license to make, have made, use, offer to sell, sell, import, and otherwise transfer the Work,

where such license applies only to those patent claims licensable by such Contributor that are necessarily infringed by their Contribution(s) alone or by combination of their Contribution(s) with the Work to which such Contribution(s) was submitted. If You institute patent litigation against any entity (including a cross-claim or counterclaim in a lawsuit) alleging that the Work or a Contribution incorporated within the Work constitutes direct or contributory patent infringement, then any patent licenses granted to You under this License for that Work shall terminate as of the date such litigation is filed.

4. Redistribution. You may reproduce and distribute copies of the Work or Derivative Works thereof in any medium, with or without modifications, and in Source or Object form, provided that You meet the following conditions:
  - (a) You must give any other recipients of the Work or Derivative Works a copy of this License; and
  - (b) You must cause any modified files to carry prominent notices stating that You changed the files; and
  - (c) You must retain, in the Source form of any Derivative Works that You distribute, all copyright, patent, trademark, and attribution notices from the Source form of the Work, excluding those notices that do not pertain to any part of the Derivative Works; and
  - (d) If the Work includes a "NOTICE" text file as part of its distribution, then any Derivative Works that You distribute must include a readable copy of the attribution notices contained within such NOTICE file, excluding those notices that do not pertain to any part of the Derivative Works, in at least one of the following places: within a NOTICE text file distributed as part of the Derivative Works; within the Source form or documentation, if provided along with the Derivative Works; or, within a display generated by the Derivative Works, if and wherever such third-party notices normally appear. The contents of the NOTICE file are for informational purposes only and do not modify the License. You may add Your own attribution notices within Derivative Works that You distribute, alongside or as an addendum to the NOTICE text from the Work, provided that such additional attribution notices cannot be construed as modifying the License.

You may add Your own copyright statement to Your modifications and may provide additional or different license terms and conditions for use, reproduction, or distribution of Your modifications, or for any such Derivative Works as a whole, provided Your use, reproduction, and distribution of the Work otherwise complies with the conditions stated in this License.

5. Submission of Contributions. Unless You explicitly state otherwise, any Contribution intentionally submitted for inclusion in the Work by You to the Licensor shall be under the terms and conditions of this License, without any additional terms or conditions. Notwithstanding the above, nothing herein shall supersede or modify the terms of any separate license agreement you may have executed with Licensor regarding such Contributions.
6. Trademarks. This License does not grant permission to use the trade names, trademarks, service marks, or product names of the Licensor, except as required for reasonable and customary use in describing the origin of the Work and reproducing the content of the NOTICE file.
7. Disclaimer of Warranty. Unless required by applicable law or agreed to in writing, Licensor provides the Work (and each Contributor provides its Contributions) on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied, including, without limitation, any warranties or conditions of TITLE, NON-INFRINGEMENT, MERCHANTABILITY, or FITNESS FOR A PARTICULAR PURPOSE. You are solely responsible for determining the

appropriateness of using or redistributing the Work and assume any risks associated with Your exercise of permissions under this License.

8. Limitation of Liability. In no event and under no legal theory, whether in tort (including negligence), contract, or otherwise, unless required by applicable law (such as deliberate and grossly negligent acts) or agreed to in writing, shall any Contributor be liable to You for damages, including any direct, indirect, special, incidental, or consequential damages of any character arising as a result of this License or out of the use or inability to use the Work (including but not limited to damages for loss of goodwill, work stoppage, computer failure or malfunction, or any and all other commercial damages or losses), even if such Contributor has been advised of the possibility of such damages.
9. Accepting Warranty or Additional Liability. While redistributing the Work or Derivative Works thereof, You may choose to offer, and charge a fee for, acceptance of support, warranty, indemnity, or other liability obligations and/or rights consistent with this License. However, in accepting such obligations, You may act only on Your own behalf and on Your sole responsibility, not on behalf of any other Contributor, and only if You agree to indemnify, defend, and hold each Contributor harmless for any liability incurred by, or claims asserted against, such Contributor by reason of your accepting any such warranty or additional liability.

END OF TERMS AND CONDITIONS

APPENDIX: How to apply the Apache License to your work.

To apply the Apache License to your work, attach the following boilerplate notice, with the fields enclosed by brackets "[]" replaced with your own identifying information. (Don't include the brackets!) The text should be enclosed in the appropriate comment syntax for the file format. We also recommend that a file or class name and description of purpose be included on the same "printed page" as the copyright notice for easier identification within third-party archives.

Copyright [yyyy] [name of copyright owner]

Licensed under the Apache License, Version 2.0 (the "License");  
you may not use this file except in compliance with the License.  
You may obtain a copy of the License at

<http://www.apache.org/licenses/LICENSE-2.0>

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

## GNU General Public License

This license covers the following software:

- RCS
- GZIP
- Ident

GNU GENERAL PUBLIC LICENSE  
Version 2, June 1991

Copyright (C) 1989, 1991 Free Software Foundation, Inc.  
675 Mass Ave, Cambridge, MA 02139, USA  
Everyone is permitted to copy and distribute verbatim copies  
of this license document, but changing it is not allowed.

Preamble



The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public License is intended to guarantee your freedom to share and change free software--to make sure the software is free for all its users. This General Public License applies to most of the Free Software Foundation's software and to any other program whose authors commit to using it. (Some other Free Software Foundation software is covered by the GNU Library General Public License instead.) You can apply it to your programs, too.

When we speak of free software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs; and that you know you can do these things.

To protect your rights, we need to make restrictions that forbid anyone to deny you these rights or to ask you to surrender the rights. These restrictions translate to certain responsibilities for you if you distribute copies of the software, or if you modify it.

For example, if you distribute copies of such a program, whether gratis or for a fee, you must give the recipients all the rights that you have. You must make sure that they, too, receive or can get the source code. And you must show them these terms so they know their rights.

We protect your rights with two steps: (1) copyright the software, and (2) offer you this license which gives you legal permission to copy, distribute and/or modify the software.

Also, for each author's protection and ours, we want to make certain that everyone understands that there is no warranty for this free software. If the software is modified by someone else and passed on, we want its recipients to know that what they have is not the original, so that any problems introduced by others will not reflect on the original authors' reputations.

Finally, any free program is threatened constantly by software patents. We wish to avoid the danger that redistributors of a free program will individually obtain patent licenses, in effect making the program proprietary. To prevent this, we have made it clear that any patent must be licensed for everyone's free use or not licensed at all.

The precise terms and conditions for copying, distribution and modification follow.

GNU GENERAL PUBLIC LICENSE  
TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

0. This License applies to any program or other work which contains a notice placed by the copyright holder saying it may be distributed under the terms of this General Public License. The "Program", below, refers to any such program or work, and a "work based on the Program" means either the Program or any derivative work under copyright law: that is to say, a work containing the Program or a portion of it, either verbatim or with modifications and/or translated into another language. (Hereinafter, translation is included without limitation in the term "modification".) Each licensee is addressed as "you".

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running the Program is not restricted, and the output from the Program

is covered only if its contents constitute a work based on the Program (independent of having been made by running the Program). Whether that is true depends on what the Program does.

1. You may copy and distribute verbatim copies of the Program's source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and give any other recipients of the Program a copy of this License along with the Program.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

2. You may modify your copy or copies of the Program or any portion of it, thus forming a work based on the Program, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:

a) You must cause the modified files to carry prominent notices stating that you changed the files and the date of any change.

b) You must cause any work that you distribute or publish, that in whole or in part contains or is derived from the Program or any part thereof, to be licensed as a whole at no charge to all third parties under the terms of this License.

c) If the modified program normally reads commands interactively when run, you must cause it, when started running for such interactive use in the most ordinary way, to print or display an announcement including an appropriate copyright notice and a notice that there is no warranty (or else, saying that you provide a warranty) and that users may redistribute the program under these conditions, and telling the user how to view a copy of this License. (Exception: if the Program itself is interactive but does not normally print such an announcement, your work based on the Program is not required to print an announcement.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Program, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Program, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Program.

In addition, mere aggregation of another work not based on the Program with the Program (or with a work based on the Program) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

3. You may copy and distribute the Program (or a work based on it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you also do one of the following:

a) Accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections

1 and 2 above on a medium customarily used for software interchange; or,

b) Accompany it with a written offer, valid for at least three years, to give any third party, for a charge no more than your cost of physically performing source distribution, a complete machine-readable copy of the corresponding source code, to be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,

c) Accompany it with the information you received as to the offer to distribute corresponding source code. (This alternative is allowed only for noncommercial distribution and only if you received the program in object code or executable form with such an offer, in accord with Subsection b above.)

The source code for a work means the preferred form of the work for making modifications to it. For an executable work, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the executable. However, as a special exception, the source code distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

If distribution of executable or object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place counts as distribution of the source code, even though third parties are not compelled to copy the source along with the object code.

4. You may not copy, modify, sublicense, or distribute the Program except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense or distribute the Program is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

5. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Program or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Program (or any work based on the Program), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Program or works based on it.

6. Each time you redistribute the Program (or any work based on the Program), the recipient automatically receives a license from the original licensor to copy, distribute or modify the Program subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties to this License.

7. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Program at all. For example, if a patent

license would not permit royalty-free redistribution of the Program by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Program.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system, which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

8. If the distribution and/or use of the Program is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Program under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.

9. The Free Software Foundation may publish revised and/or new versions of the General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Program specifies a version number of this License which applies to it and "any later version", you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Program does not specify a version number of this License, you may choose any version ever published by the Free Software Foundation.

10. If you wish to incorporate parts of the Program into other free programs whose distribution conditions are different, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

#### NO WARRANTY

11. BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

12. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

END OF TERMS AND CONDITIONS

Appendix: How to Apply These Terms to Your New Programs

If you develop a new program, and you want it to be of the greatest possible use to the public, the best way to achieve this is to make it free software which everyone can redistribute and change under these terms.

To do so, attach the following notices to the program. It is safest to attach them to the start of each source file to most effectively convey the exclusion of warranty; and each file should have at least the "copyright" line and a pointer to where the full notice is found.

< signature of Ty Coon 1 April 1989  
Ty Coon, President of Vice

This General Public License does not permit incorporating your program into proprietary programs. If your program is a subroutine library, you may consider it more useful to permit linking proprietary applications with the library. If this is what you want to do, use the GNU Library General Public License instead of this License.

## TCL License

This software is copyrighted by the Regents of the University of California, Sun Microsystems, Inc., Scriptics Corporation, and other parties. The following terms apply to all files associated with the software unless explicitly disclaimed in individual files.

The authors hereby grant permission to use, copy, modify, distribute, and license this software and its documentation for any purpose, provided that existing copyright notices are retained in all copies and that this notice is included verbatim in any distributions. No written agreement, license, or royalty fee is required for any of the authorized uses.

Modifications to this software may be copyrighted by their authors and need not follow the licensing terms described here, provided that the new terms are clearly indicated on the first page of each file where they apply.

IN NO EVENT SHALL THE AUTHORS OR DISTRIBUTORS BE LIABLE TO ANY PARTY FOR DIRECT, INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OF THIS SOFTWARE, ITS DOCUMENTATION, OR ANY DERIVATIVES THEREOF, EVEN IF THE AUTHORS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

THE AUTHORS AND DISTRIBUTORS SPECIFICALLY DISCLAIM ANY WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, AND NON-INFRINGEMENT. THIS SOFTWARE IS PROVIDED ON AN "AS IS" BASIS, AND THE AUTHORS AND DISTRIBUTORS HAVE NO OBLIGATION TO PROVIDE MAINTENANCE, SUPPORT, UPDATES, ENHANCEMENTS, OR MODIFICATIONS.

GOVERNMENT USE: If you are acquiring this software on behalf of

the U.S. government, the Government shall have only "Restricted Rights" in the software and related documentation as defined in the Federal Acquisition Regulations (FARs) in Clause 52.227.19 (c) (2). If you are acquiring the software on behalf of the Department of Defense, the software shall be classified as "Commercial Computer Software" and the Government shall have only "Restricted Rights" as defined in Clause 252.227-7013 (c) (1) of DFARs. Notwithstanding the foregoing, the authors grant the U.S. Government and others acting in its behalf permission to use and distribute the software in accordance with the terms specified in this license.

## Perl License

### Preamble

The intent of this document is to state the conditions under which a Package may be copied, such that the Copyright Holder maintains some semblance of artistic control over the development of the package, while giving the users of the package the right to use and distribute the Package in a more-or-less customary fashion, plus the right to make reasonable modifications.

### Definitions

"Package" refers to the collection of files distributed by the Copyright Holder, and derivatives of that collection of files created through textual modification.

"Standard Version" refers to such a Package if it has not been modified, or has been modified in accordance with the wishes of the Copyright Holder as specified below.

"Copyright Holder" is whoever is named in the copyright or copyrights for the package.

"You" is you, if you're thinking about copying or distributing this Package.

"Reasonable copying fee" is whatever you can justify on the basis of media cost, duplication charges, time of people involved, and so on. (You will not be required to justify it to the Copyright Holder, but only to the computing community at large as a market that must bear the fee.)

"Freely Available" means that no fee is charged for the item itself, though there may be fees involved in handling the item. It also means that recipients of the item may redistribute it under the same conditions they received it.

1. You may make and give away verbatim copies of the source form of the Standard Version of this Package without restriction, provided that you duplicate all of the original copyright notices and associated disclaimers.
2. You may apply bug fixes, portability fixes and other modifications derived from the Public Domain or from the Copyright Holder. A Package modified in such a way shall still be considered the Standard Version.
3. You may otherwise modify your copy of this Package in any way, provided that you insert a prominent notice in each changed file stating how and when you changed that file, and provided that you do at least ONE of the following:
  - a. place your modifications in the Public Domain or otherwise make them Freely Available, such as by posting said modifications to Usenet or an equivalent medium, or placing the modifications on a major archive site such as uunet.uu.net, or by allowing the Copyright Holder to include your modifications in the Standard Version of the Package.
  - b. use the modified Package only within your corporation or organization.
  - c. rename any non-standard executables so the names do not conflict with standard executables, which must also be provided, and provide a separate manual page for each non-standard executable that clearly documents how it differs from the Standard Version.
  - d. make other distribution arrangements with the Copyright Holder.
4. You may distribute the programs of this Package in object code or executable form, provided that you do at least ONE of the following:
  - a. distribute a Standard Version of the executables and library files, together with instructions (in the manual page or equivalent) on where to get the Standard Version.
  - b. accompany the distribution with the machine-readable source of the Package with your modifications.

c. give non-standard executables non-standard names, and clearly document the differences in manual pages (or equivalent), together with instructions on where to get the Standard Version.

d. make other distribution arrangements with the Copyright Holder.

5. You may charge a reasonable copying fee for any distribution of this Package. You may charge any fee you choose for support of this Package. You may not charge a fee for this Package itself. However, you may distribute this Package in aggregate with other (possibly commercial) programs as part of a larger (possibly commercial) software distribution provided that you do not advertise this Package as a product of your own. You may embed this Package's interpreter within an executable of yours (by linking); this shall be construed as a mere form of aggregation, provided that the complete Standard Version of the interpreter is so embedded.

6. The scripts and library files supplied as input to or produced as output from the programs of this Package do not automatically fall under the copyright of this Package, but belong to whomever generated them, and may be sold commercially, and may be aggregated with this Package. If such scripts or library files are aggregated with this Package via the so-called "undump" or "unexec" methods of producing a binary executable image, then distribution of such an image shall neither be construed as a distribution of this Package nor shall it fall under the restrictions of Paragraphs 3 and 4, provided that you do not represent such an executable image as a Standard Version of this Package.

7. C subroutines (or comparably compiled subroutines in other languages) supplied by you and linked into this Package in order to emulate subroutines and variables of the language defined by this Package shall not be considered part of this Package, but are the equivalent of input as in Paragraph 6, provided these subroutines do not change the language in any way that would cause it to fail the regression tests for the language.

8. Aggregation of this Package with a commercial distribution is always permitted provided that the use of this Package is embedded; that is, when no overt attempt is made to make this Package's interfaces visible to the end user of the commercial distribution. Such use shall not be construed as a distribution of this Package.

9. The name of the Copyright Holder may not be used to endorse or promote products derived from this software without specific prior written permission.

10. THIS PACKAGE IS PROVIDED "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.

## Zlib License

License

```
/* zlib.h -- interface of the 'zlib' general purpose compression library
version 1.2.1, November 17th, 2003
```

Copyright (C) 1995-2003 Jean-loup Gailly and Mark Adler

This software is provided 'as-is', without any express or implied warranty. In no event will the authors be held liable for any damages arising from the use of this software.

Permission is granted to anyone to use this software for any purpose, including commercial applications, and to alter it and redistribute it freely, subject to the following restrictions:

1. The origin of this software must not be misrepresented; you must not claim that you wrote the original software. If you use this software in a product, an acknowledgment in the product documentation would be appreciated but is not required.

2. Altered source versions must be plainly marked as such, and must not be misrepresented as being the original software.

3. This notice may not be removed or altered from any source distribution.

Jean-loup Gailly [jloup@gzip.org](mailto:jloup@gzip.org)  
Mark Adler [madler@alumni.caltech.edu](mailto:madler@alumni.caltech.edu)