Tru64 UNIX

Release Notes for Versions 5.1B-1 and 5.1B-2

Part Number: AA-RVGEB-TE

August 2004

Product Version: Tru64 UNIX Versions 5.1B-1 and 5.1B-2

This manual provides information on new and changed features for the HP Tru64 UNIX operating system. It also provides information on restrictions to the software and documentation.

You can also view the *Technical Update* for Version 5.1B or higher for any additional information not included in these notes. You can access the *Technical Update* from the following URL:

http://h30097.www3.hp.com/docs/updates/V51B/TITLE.HTM

Hewlett-Packard Company
Palo Alto, California

Contents

4 Software Notes for Tru64 UNIX Version 5.1B-1

Index

About This Manual

This manual contains release notes for the HP Tru64 UNIX Version 5.1B-1 and Version 5.1B-2 operating system software.

This manual also describes significant new and changed features in these versions of the Tru64 UNIX operating system, and lists features and interfaces scheduled for retirement in future releases.

Audience

These release notes are for the person who installs the product and for anyone using the product following installation.

Organization

This manual is organized as follows:

Chapter 1	Contains an overview of new and changed features in Version 5.1B-2 of the operating system software
Chapter 2	Contains information on restrictions to the software and documentation in Version 5.1B-2 of the operating system software
Chapter 3	Contains an overview of new and changed features in Version 5.1B-1 of the operating system software
Chapter 4	Contains information on restrictions to the software and documentation in Version 5.1B-1 of the operating system software

Related Documents

You will find it helpful to have the following documentation available during the installation of this product:

- The hardware documentation for your system
- The *Installation Guide*
- The *Installation Guide — Advanced Topics*
- The online reference pages
- The HTML files provided on the Software Documentation CD–ROM, especially *New and Changed Features from Previous Releases*

You can also view the *Technical Update* for Version 5.1B or higher for any additional information not included in these notes. You can access the *Technical Update* from the following URL:

http://h30097.www3.hp.com/docs/updates/V51B/TITLE.HTM

Reader's Comments

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

You can send your comments in the following ways:

- Fax: 603-884-0120 Attn: UBPG Publications, ZKO3-3/Y32
- Internet electronic mail: `readers_comment@zk3.dec.com`

  A Reader's Comment form is located on your system in the following location:

  `/usr/doc/readers_comment.txt`

Please include the following information along with your comments:

- The full title of the manual and the order number, if the manual has one. (When provided, the order number appears on the title page of printed and PDF versions of a manual.)
- The section numbers and page numbers of the information on which you are commenting.
- The version of Tru64 UNIX that you are using.
- If known, the type of processor that is running the Tru64 UNIX software.

The Tru64 UNIX Publications group cannot respond to system problems or technical support inquiries. Please address technical questions to your local system vendor or to the appropriate HP technical support office. Information

provided with the software media explains how to send problem reports to HP.

Conventions

The following conventions are used in this manual:

`%` `$`	A percent sign represents the C shell system prompt. A dollar sign represents the system prompt for the Bourne, Korn, and POSIX shells.
`#`	A number sign represents the superuser prompt.
`% cat`	Boldface type in interactive examples indicates typed user input.
`file`	Italic (slanted) type indicates variable values, placeholders, and function argument names.
`[ | ]` `{ | }`	In syntax definitions, brackets indicate items that are optional and braces indicate items that are required. Vertical bars separating items inside brackets or braces indicate that you choose one item from among those listed.
`...`	In syntax definitions, a horizontal ellipsis indicates that the preceding item can be repeated one or more times.
`cat(1)`	A cross-reference to a reference page includes the appropriate section number in parentheses. For example, `cat(1)` indicates that you can find information on the `cat` command in Section 1 of the reference pages.
Ctrl/*x*	This symbol indicates that you hold down the first named key while pressing the key or mouse button that follows the slash. In examples, this key combination is enclosed in a box (for example, Ctrl/C).

1

New and Changed Features for Tru64 UNIX Version 5.1B-2

This chapter describes new features that are available with the Version 5.1B-2 release of the operating system. It also lists new hardware that is supported.

1.1 New Hardware Support

The following new hardware support has been added in this release:

- This release provides support for the upgraded EV7 Alpha chip operating at 1.3 GHz.
- HP StorageWorks DAT 40x6 Tape Autoloader: The HP StorageWorks DAT 40x6 tape autoloader is a dependable, entry-level solution for small-to-medium server storage and enterprise network backup needs.
- HP StorageWorks DLT VS80: The StorageWorks DLT VS80 tape drive provides affordable 80-GB backup to IT managers with midrange servers.
- FCA2684 and FCA2684DC: This adapter is a PCI-X 2–GB 64-bit/133MHz single and dual port Fibre Channel Host Bus Adapter.

1.2 New Functionality

The following section discusses new functionality for Tru64 UNIX Version 5.1B-2.

1.2.1 Inclusive Patch Kits

The 5.1B-2/PK4 patch kit marks a new way of delivering Tru64 UNIX patches. If you have installed previous Tru64 UNIX patch kits, you will see the following differences when you install this kit:

- 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 sure 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 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 25001.00 - SP04 OSFACCT540

      * Patch 25002.00 - SP04 OSFADVFS540 (SSRT2275)

      * Patch 25003.00 - 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 the 5.1B-2/PK4 patch kit, the `OSF...540` identifies the subset the patch is included in, and the `SSRT2275` indicates a type of security patch.

  As with previous kits, you can find a brief overview of all the patches (listed by patch number) in the kit's *Patch Summary and Release Notes*.

- 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

  Beginning with the 5.1B-2/PK4 patch kit, patch kits will include any 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.

Except for the installation and removal processes, the functions provided by the `dupatch` utility generally work the same with inclusive patch kits as they do in old-style patch kits. For example, the Patch Tracking and Baselining menus remain the same and work the same in the 5.1B-2/PK4 patch kit as they do in the old-style patch kits.

The *Patch Kit Installation Instructions* manual provides information for installing inclusive patch kits and the old-style kits. Where the processes differ, each process is explained.

You can install the 5.1B-2/PK4 patch kit on any system running Version 5.1B or Version 5.1B-1.

1.2.2 Unified Buffer Cache Scaling

This release introduces Unified Buffer Cache Scaling, which improves large SMP and NUMA system performance. MSI interrupts entails implementing the PCI-X specification to allow the device to use >1 MSI. With EV7 and EV7z being PCI-X systems with MSI-capable devices and drivers, they are sure to see better performance from reduced locks, especially in large configurations.

1.3 TruCluster Server Improvements

The following sections describe TruCluster Server improvements for Tru64 UNIX Version 5.1B-2/PK4.

1.3.1 Faster Boot Times for TruCluster Server Environments Using Logical Storage Manager

This release contains enhancements that provide a significant reduction in the time required to boot TruCluster Server configurations using LSM.

1.3.2 Improved Memory Channel Performance During Peak System Utilization

The Memory Channel subsystem has been enhanced to support parallel Memory Channel input processing. This enhancement increases the Memory Channel performance, and provides resistance to CPU starvation.

1.3.3 Additional TruCluster Enhancements

The following TruCluster Server enhancements are new for this release

- Sticky Cluster Connections: This is an attribute for a cluster alias that will direct all client connections to the same cluster node, providing a stickiness of a servicing cluster member to a specific client/port.
- Improved Cluster Serviceability: Improves the troubleshooting efficiency by providing a mechanism to retrieve runtime information from the cluster alias daemon (`aliasd`).

- Improved Cluster Manageability: TruCluster Server now supports the `sysconfig` command across the entire clusters, in much the same way that it currently runs on a single system.

  In a TruCluster Server environment, the `sysconfig` command uses the cluster interconnect to send requests to reconfigure, query attributes, and query subsystem states of kernel subsystems on different cluster members. The `sysconfig` command receives output from these commands across the cluster interconnect. Using the cluster interconnect for these commands allows querying or modification attributes on members that are hung or on members that do not have an external interface between cluster members.

  The cluster interconnect is not used for the `sysconfig configure` and `unconfigure` commands.

1.4 Networking Improvements

The following sections discuss the networking improvements for Tru64 UNIX Version 5.1B-2.

1.4.1 Mobile IPv6 Update

This release provides an update to the current Mobile IPv6 code to bring it in line with the latest Networking RFCs.

1.4.2 IPv6 Advanced API Update

The IPv6 advanced API has been updated to conform to RFC3542.

1.4.3 pfilt_loopback and pfilt_physaddr

Two packetfilter tunable variables, `pfilt_loopback` and `pfilt_physaddr`, were previously tunable only by using `dbx`. These variables are now tunable using `sysconfig` and `/etc/sysconfigtab`. See Section 3.4 for more information about these variables.

1.5 Worldwide Language Support Improvements

The following sections describe the Worldwide Language Support improvements for Tru64 UNIX Version 5.1B-2.

1.5.1 Updated Localized Messages

Localized messages have been updated to match English messages that have changed since the release of Tru64 UNIX Version 5.1B.

1.5.2 Improved Asian TTY Subsystem

The Asian TTY subsystem has been improved to reduce the risk of kernel crashes under heavy workloads.

1.5.3 Enhanced Chinese Support

The Qu-Wei input method invoked from `dxhanziim` or `dxim` now produces a correct character from a Unicode value.

Characters are now drawn with correct width in the zh_CN.GB18030 locale.

1.5.4 Improved Iconv Converters

Iconv converters for Japanese mainframe codesets (ibmkanji/JEF/KEIS) have been improved to reduce the number of incorrect result.

1.6 New and Updated Associated Products

Several software products provided on the Associated Products CD–ROMs have been updated for this release. The updated products are listed in the following sections.

For more information on the CD–ROM contents, see the *HP Tru64 UNIX Version 5.1B-2 CD-ROMs* card contained in the media kit.

1.6.1 Advanced Server for Tru64 UNIX

The Advanced Server for Tru64 UNIX software has been updated to Version 5.1B-2. This update provides enhancements and corrections for problems found in the ASU Version 5.1B ECO2 software, and in earlier versions of the ASU software.

Some of the enhancements and corrections in the ASU Version 5.1B-2 kit:

- The ASU server can now be a Backup Domain Controller to a Windows 2003 mixed mode domain.
- The ASU server now supports the nondefault cluster alias. The ASU server supports only one cluster alias at a time, which can be either the default alias or a nondefault alias.
- The FileChangeNotify registry parameter has been added to enable the ASU server to notify clients of a file change from other clients.
- The StoreAttributesAsMetadata registry parameter has been added to enable storing of DOS attributes and file creation time in AdvFS metadata. This allows the use of DOS attributes irrespective of the UseUnixGroups registry parameter value, and preserves the file creation

time when modifying a file. The default value is 0 (do not store DOS attributes and file creation time in AdvFS metadata).

- Three new ASU commands `chattr`, `lsattr`, and `rmattr`, are added to manage DOS attributes. See the associated reference pages for a description of each command.
- A new command, `prcheck`, is available to check and enumerate ASU printer entries, which includes the ASU printer share entries, printer printcap entries, printer registry entries, and printer spool directory entries.
- A new client tool, `pccheck`, is available to collect and display the diagnostic information such as user, share, network statistics, browse list, and connectivity to PDC/BDC.

1.6.2 Tru64 UNIX to HP-UX Software Transition Kit

The Tru64 UNIX to HP-UX Software Transition Kit has been updated to Version 2.2. The Tru64 UNIX to HP-UX Software Transition Kit has been updated to include a broader coverage of commands, libraries and programming languages. Tru64 UNIX developers can now use the STK to scan Fortran source code files in addition to C and C++, makefiles and shell scripts.

1.6.3 XEmacs

XEmacs has been updated to Version 21.4 with this release.

1.6.4 Secure Web Server

The Secure Web Server has been updated to Version 6.3.0 with this release.

1.6.5 Perl

Perl has been updated to Version 5.8.4 with this release.

1.6.6 Extended System V Functionality

The Extended System V Functionality has been updated to Version 2.1. This update includes new and modified commands and APIs, and updated reference pages.

1.6.7 UniCensus

UniCensus has been upgraded to Version 5.0.5 with this release.

1.6.8 Visual Threads

Visual Threads has been updated to Version 2.4–003 with this release.

1.6.9 Web Based Enterprise Service

The Web Based Enterprise Service (WEBES) has been updated to Version 4.3.3 with this release.

1.6.10 Collect GUI

The Collect GUI has been updated to version 2.0.6 with this release.

1.6.11 OpenLDAP Utilities

LDAP Utilities is updated to include the latest OpenLDAP tools from version 2.1.25, minor bug fixes, and general improvements.

1.6.12 OpenLDAP Directory Server

OpenLDAP has been updated to version 2.1.25. This update includes bug fixes and general improvements.

1.6.13 Mozilla

Mozilla has been upgraded to Version 1.6. For a description of the new features provided in this version, see the Release Notes at the following URL:

http://www.mozilla.org/releases/mozilla1.6/README.html#new

1.6.14 Legato NetWorker

Legato NetWorker for HP Tru65 UNIX has been updated to Version 7.1. This update includes a number of new and advanced features.

2

Software Notes for Tru64 UNIX Version 5.1B-2

2.1 Installation Notes

The following sections describe installation notes for Tru64 UNIX Version 5.1B-2.

2.1.1 Possible Error Seen During Kit Installation

You may see the following messages during the installation process:

```
1200600:/sbin/hwmgr: /sbin/loader: Error:
         libpthread.so: symbol "_callback_rmutex" unresolved
1200600:/sbin/hwmgr: /sbin/loader: Fatal Error:
        Load of "/sbin/hwmgr" failed: Unresolved symbol name
```

These messages are generated as a result of the order in which components are installed and can be ignored.

2.1.2 Possible Error Seen After Kit Installation

The following problems have been known to occur after the 5.1B-2/PK4 patch kit 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/
```

2.1.3 Message Seen During Reboot Can Be Ignored

The following error message will be displayed after you reboot your system the first time after installing the 5.1B-2/PK4 patch kit:

```
AllowCshrcSourcingWithSubsystems is not valid
ForcePTTYAllocation is not valid
IdentityFile is not valid
AuthorizationFile is not valid
```

These messages are caused by a new version of SSH included in the 5.1B-2/PK4 patch kit. They do not pose a problem and can be ignored.

2.1.4 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.

2.1.5 Required Actions When Uninstalling the 5.1B-2/PK4 Patch Kit

The following sections describe actions you have to take if you decided to uninstall the 5.1B-2/PK4 patch kit. Read each section before running the patch deletion procedure.

2.1.5.1 Script Required to Reverse Version Switch

If you enabled version switches as described in Section 2.1.4, you must run the `/usr/sbin/versw_enable_delete` script before attempting to remove the 5.1B-2/PK4 patch kit. 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 the 5.1B-2/PK4 patch kit (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 the 5.1B-2/PK4 patch kit. Because the no-roll procedure automatically reboots the system after deleting the patches, you would not be able to perform this step as required.

2.1.5.2 Changes to System May Need to Be Reversed

If you made the following changes to your system after installing the 5.1B-2/PK4 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 the kit might not recognize the new devices or may not provide the necessary support for them.
- If you added new cluster members, the new members will not have an older state to revert to if you attempt to uninstall the kit.

To uninstall the 5.1B-2/PK4 patch kit, do the following:

1. Remove all new hardware and new cluster members that you added after installing the 5.1B-2/PK4 patch kit.
2. Run `dupatch` to uninstall the patch kit.
3. Verify that the patch kit was successfully uninstalled.

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

2.1.5.3 Script Required When Returning to Prepatched System

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

This situation occurs if the 5.1B-2/PK4 patch kit is the only Tru64 UNIX patch kit installed on your 5.1B system.

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 the system.

If you also need to reverse the version switch as described in Section 2.1.5.1, run the `/etc/dn_fix_dat.sh` script after step 5 in that process.

Note

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

2.1.6 Do Not Use Command Line to Remove the 5.1B-2/PK4 Patch Kit

If you need to uninstall the 5.1B-2/PK4 patch kit, we recommend that you use `dupatch` in interactive mode rather than from the command line. Because you must remove all patches installed with this kit, you would have to specify each of this kit's patches on the command line. With more than 60 patches in this kit, removing them using the command line would be a difficult, time-consuming task.

Although you can use the `-delete -patch all` option on the command line, doing so removes all `setld`-installed patches from your system, including CSPs and the patches from previous patch kits. Therefore, use this option only if patches from this patch kit are the only ones installed on your system. Use the `-track` option to determine if your system contains any other patches.

By using `dupatch` in interactive mode, selecting `delete` from the menu causes `dupatch` to display all of the patches on your system. You can then easily select only the patches in the 5.1B-2/PK4 patch kit. See the examples appendix of the *Patch Kit Installation Instructions* for a complete example of removing this patch kit using `dupatch` interactively.

2.1.7 Additional Steps Required for HP Insight Management Agents Kit

Under certain conditions, you will be prevented from installing the 5.1B-2/PK4 patch kit if you are running HP Insight Management Agents kit CPQIM310 or higher or had a version of the kit previously installed. Those conditions are as follows:

- Your system contains an earlier patch kit than the 5.1B-2/PK4 patch kit and the Insight Management Agents kit.

In this case, upgrading to the 5.1B-2/PK4 patch kit gives the following error message:

```
Patch 25020.00 - SP04 OSFCLINET540 (SSRT3653 SSRT2384 SSRT2275 ...)
./sbin/init.d/snmpd: its origin can not be identified.

This patch will not be installed.
```

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

  In this case, upgrading to the 5.1B-2/PK4 patch kit gives the following error message:

```
 Patch 25020.00 - SP04 OSFCLINET540 (SSRT3653 SSRT2384 SSRT2275...)
 ./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 run the `dupatch` baseline process before installing the 5.1B-2/PK4 patch kit. The following steps will guide you through the process:

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

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

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

2. Run the `dupatch` utility and select Option 5, `Patch Baseline Analysis/Adjustment`. See the *Patch Kit Installation Instructions* for detailed 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 25020.00 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 the 5.1B-2/PK4 patch kit.

5. After the system is running with the 5.1B-2/PK4 patch kit 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 /temp/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 the 5.1B-2/PK4 patch kit (step 4) and stop the snmpd and insightd daemons (step 5) as follows (the XXX represents the revision, such as 310):

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 &
```

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 the 5.1B-2/PK4 patch kit, which includes Patch 25020, 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. You do this by restoring the backup version you created in step 1 of the workaround procedure, or by modifying the replacement script as described in this section.

2.1.8 Required Action when Installing Extended System V Functionality and Tru64 UNIX Worldwide Language Support Subsets from the APCDs

To successfully install the Extended System V Functionality and the Tru64 UNIX Worldwide Language Support subsets, on the same system, the following procedure must be performed:

1. Install the Tru64 UNIX Worldwide Language Support subsets first.
2. Install the Extended System V Functionality subsets.
3. Copy the following files:

```
# cp /usr/i18n/lib/nls/loc/iconvTable/* /usr/lib/nls/loc/iconvTable
```

2.2 Max_LSM_IO_PERFORMANCE Restriction with root or cluster_root File Systems Under LSM

Setting the sysconfigtab variable Max_LSM_IO_PERFORMANCE = 1, while the root or cluster root domain is under LSM control, will cause a system hang during the boot process.

This feature was added to improve performance on multiprocessor systems when running heavy I/O loads. To enable this feature, the user needs to set the Max_LSM_IO_PERFORMANCE sysconfig variable.

See sys_attrs_lsm(5) for more information.

2.3 LSM Merge Problem During Update Installation on a Standalone System with Mirrored Root File System

During an update installation, the installupdate command does not merge the LSM entries into /etc/inittab, which prevents LSM from automatically starting up on a system reboot until later in the boot process when /sbin/bcheckrc is run.

In most cases, this does not cause a problem and is easily fixed by running the /usr/sbin/volinstall command, which updates the /etc/inittab file with the LSM entries.

However, if the root file system is mirrored (for example via volencap and volrootmir) and LSM does not start until multiuser mode (via bcheckrc), the system may write to only one of the root file system mirrors, causing the other mirrored plex (the one not associated with the boot disk) to contain inconsistent data.

To avoid this, remove the mirrored plex not associated with the boot disk before performing the update installation. After the installation is completed, run /usr/sbin/volinstall to update the /etc/inittab file with the missing LSM entries and run /usr/sbin/volrootmir if you want to mirror the root file system again.

The following steps show how to accomplish this. In the examples, the boot disk is dsk6 and the mirrored disk is dsk5.

1. Before running the installupdate command, do the following:

   a. Check to see if the boot disk is mirrored, for example rootvol-01 and rootvol-02:

      ```
      # /sbin/volprint -g rootdg
      ```

   b. If the root file system is mirrored, remove the second plexes associated with the boot disk to avoid the system writing to only one of the root plexes during the installation.

```
# /sbin/volplex -o rm dis rootvol-02 swapvol-02 vol-dsk6g-02
```

c. Remove the associated disks from LSM control. This is necessary if the root file system is to be mirrored again to this disk:

```
# /sbin/voldg rmdisk root02 swap02 dsk5g-AdvFS dsk5e

# /sbin/voldisk rm dsk5a dsk5b dsk5g dsk5e
```

2. Run the update installation. For example:

```
# /sbin/installupdate /dev/disk/cdrom0c
```

3. When the installation is completed, do the following:

a. Add LSM entries into the `/etc/inittab` file to allow LSM to start automatically on system boot:

```
# /usr/sbin/volinstall
```

b. If you want to mirror the root file system, you can do so now:

```
# /usr/sbin/volrootmir dsk5
```

2.4 Base Operating System Notes

The following sections describe the base operating system notes for Tru64 UNIX Version 5.1B-2.

2.4.1 Mozilla Does Not Accept Keyboard Input When the NumLock Key is Enabled

When the `NumLock` is on (on a PC keyboard), Mozilla does not accept any keyboard input. You must turn off the `NumLock` key when you use Mozilla.

2.4.2 Messages Displayed During XEmacs Startup

The following messages are displayed when XEmacs Version 21.4 is started:

```
Loading xlib-math...
Loading xlib-math...done
Loading xwem-compat...
Loading xwem-compat...done
```

These messages can be ignored.

2.4.3 Input Methods Do Work With Motif Version 2.1

The `/usr/opt/motif2.1/README` file states that input methods, which currently work with Motif 1.2, do not work with Motif 2.1. This statement is incorrect. In most cases, the input methods work.

2.4.4 Potential False Temperature Error Condition on AlphaServer DS10, DS10L, and TS10 Systems

A sensor error can potentially indicate a false over-temperature condition on AlphaServer DS10, DS10L, and TS10 systems. While the sensor accurately reports the temperature, it falsely reports the high temperature threshold. This false reporting of the threshold value can cause the Insight Management SMP Agents to forward false traps, and potentially result in the Tru64 UNIX environmental monitoring daemon (envmond(8)) initiating the shutdown process. However, this false reporting condition is temporary and the high temperature threshold soon returns to normal, at which point the shutdown is cancelled.

To work around this problem, use the envconfig utility to manually set the high temperature threshold to its current default value. While this does not change the value, the act of setting it manually forces envmond to use the manually applied value. Once manually set, the ENVMON_HIGH_THRESH variable persists in the /etc/rc.config database. This will permanently work around the issue for current and future envmond sessions. To manually set the variable, follow these steps:

1. Determine the the value of ENVMON_HIGH_THRESH:

```
# sysconfig -q envmon | grep thresh
# env_high_temp_thresh = 60
```

2. Manually set the threshold value using the envconfig utility:

```
# envconfig -c ENVMON_HIGH_THRESH=60
```

2.4.5 Configuring IPsec

The*SSRT3674 - HP Tru64 UNIX IPsec/IKE Potential Security Vulnerability* HP Security Bulletin identifies a potential security vulnerability in the HP Tru64 UNIX operating system using IPsec/IKE (Internet Key Exchange) with Certificates. The potential vulnerability may be remotely exploitable, resulting in unauthorized privileged access.

HP has corrected this potential vulnerability by releasing a fix in the 5.1B-2/PK4 patch kit and the following updated documentation. The updated functionality allowing for restriction of remote identities (IDs) is outlined in step 12 for securely configuring a host or gateway. This updated functionality is only available after installing the 5.1B-2/PK4 patch kit.

Use the SysMan Menu application of the Common Desktop Environment (CDE) Application Manager to configure IPsec. This section describes how to configure your system as either an IPsec host or a secure gateway.

2.4.5.1 Configuring a Host

To configure IPsec on a host, follow these steps:

1. From the SysMan Menu, select Networking→Additional Network Services→Configure Internet Protocol Security (IPsec) to display the IPsec main window.

   Alternatively, enter the following command on the command line:

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

   If you are configuring IPsec for the first time, an informational dialog box is displayed that tells you to define secure connections before enabling IPsec. If you enable IPsec without defining secure connections, all packets into and out of the system are discarded; no traffic will flow. Select OK.

   The IPsec main window displays configured secure connections and configured public-key certificates.
2. Select Enable IP Security (IPsec) at the top of the window.
3. Select Add. The Add/Modify a Secure Connection dialog box is displayed.
4. Enter a connection name.
5. Select Add to add a remote IP address selector. The Add/Modify Selector dialog box is displayed. Do the following:

   a. Select a selector type.

   b. Do one of the following:

      - If you are communicating with a single host, enter the IP address.
      - If you are communicating with a secure gateway, enter the subnet address.
      - If you are communicating with a range of addresses, enter the first address.

   c. For an IP subnet, enter the size of the subnet mask.

   d. For a range of addresses, enter the last address.

   e. Select an upper layer protocol to match. By default, all protocols are selected.

   f. If you want to restrict the selector to a specific port number, enter a port number to match. By default, all port numbers are selected.

   g. Select OK to accept the data and close the Add/Modify Selector dialog box. If you are finished adding remote and local addresses, go to step 7.

6. Select Add to add a local IP address selector. Go to step 5a.
7. Select an action to apply to the packets matching the selectors. The default is to apply IPsec protection.
8. Select Next to accept the data and close the Add/Modify a Secure Connection dialog box. The Add/Modify Connection: IPsec Proposal dialog box is displayed. Do the following:

    a. Select an IPsec proposal from the proposal list.

    b. If you are communicating with a secure gateway, specify the IP address of the secure gateway (remote) and your system's IP address (local).

    c. Specify if you will use IKE to obtain keys or use manual configuration. Select Next to accept the data and close the Add/Modify Connection: IPsec Proposal dialog box.

    If you selected manual configuration and have created a custom proposal list with only one proposal, the Add/Modify Connection: Manual Keys dialog box displays. Go to step 9. If you selected the IKE protocol, the Add/Modify Connection: IKE Proposal dialog box displays. Go to step 11.

9. Select Add to add a manual key and display the Modify Keys: Add/Modify IPsec Key dialog box. Do the following:

    a. Enter the key name.

    b. Enter the Security Parameter Index (SPI).

    c. Enter keys for the algorithms that are required by the proposals you chose. Select OK to accept the data and close the Modify Keys: Add/Modify IPsec Key dialog box.

10. Select whether you want to apply the key(s) to inbound packets or outbound packets, or both. If you want to specify additional keys, go to step 9. If you are finished specifying manual keys, go to step 20.
11. Select an IKE proposal from the proposal list.
12. Select Add to restrict access to the connection and display the Add/Modify Remote IDs dialog box. Do the following:

    a. Select a remote identity type.

    b. Enter an identity string, usually your IP address, domain name, or e-mail address.

    c. Select OK to accept the data and close the Add/Modify Remote IDs dialog box.

Note

A remote identity (ID) is one that is allowed to use this connection. Identities are values that are either specified in a certificate by the Subject Alternate Name or that you enter when specifying a preshared key. This step is optional. However, if you do not specify a remote identity or identities, other systems might have unauthorized access to your system.

13. If you want to specify additional remote identities, go to step 12. If you are finished specifying remote identities, select Next to accept the data, close the Add/Modify Connection: IKE Proposal dialog box, and display the Add/Modify Connection: IKE Authentication dialog box.
14. Select whether you want to authenticate IKE exchanges with a public-key certificate or a preshared key.
15. If you selected public-key certificate, select Add to add an IKE certificate. The Add/Modify Certificates dialog box is displayed. Do the following:
    a. Enter a certificate name, select a certificate encoding method, and enter the local path to the certificate file.
    b. If the certificate authenticates your system, select the encoding method and enter the local path to the private key file.
    c. If the certificate is trusted to sign other certificates, select CA Certificate. Otherwise, go to step f.
    d. If a Certificate Revocation List (CRL) is not available, select No Certificate Revocation List (CRL) Available. Go to step f.
    e. Select an encoding method for the CRL and enter a local path to the CRL file.
    f. Select OK to accept the data and close the Add/Modify Certificates dialog box.
16. Select a certificate for the IKE exchange. Go to step 19.
17. If you selected a preshared key, select Add an IKE preshared key. The Add/Modify IKE Keys dialog box is displayed. Do the following:
    a. Enter a key name and key value.
    b. Select a local identity type.
    c. Enter an identity string, usually your IP address or domain name.
    d. Select OK to accept the data and close the Add/Modify IKE Keys dialog box.

18. Select a preshared key for the IKE exchange.
19. Select Next to close the Add/Modify Connection: IKE Authentication dialog box and display the Add/Modify Connection: Optional IKE Parameters dialog box. Do the following:

    a. Select any optional parameters.
    b. Select an IKE group number for initial Diffie-Hellman exchanges, if it is different from the IKE proposals.
    c. If you are using Perfect Forward Secrecy (PFS), select a group number future for Diffie-Hellman exchanges.
    d. Select a default lifetime if the proposal does not specify a lifetime.
    e. Select Finish to accept the data and close the Add/Modify Connection: Optional IKE Parameters dialog box.

20. An informational dialog box is displayed that tells you the connection has been created. Select OK to close this dialog box.
21. If you need to specify additional public-key certificates, select Add in the Public-Key Certificates field to display an Add/Modify Certificates dialog box into which you can enter information for the certificate. Do the following:

    a. Enter the certificate name, select a certificate encoding method, and enter a local path to the certificate file.
    b. If the certificate authenticates your system, select a private key encoding method and enter a local path to the private key file.
    c. If the certificate is trusted to sign other certificates, select CA Certificate. Otherwise, go to step f.
    d. If a Certificate Revocation List (CRL) is not available, select No Certificate Revocation List (CRL) Available. Go to step f.
    e. Select an encoding method for the CRL and enter a local path to the CRL file.
    f. Select OK to accept the data and close the Add/Modify Certificates dialog box.

22. Select OK in the IPsec main window to save the configuration information. Whether or not IPsec is already running on your system, the Restart IPsec? dialog box is displayed. If you want to start or restart IPsec, select OK; otherwise, select No. If you select No, you must reboot the system to start or restart IPsec.

See the *Network Administration: Connections* manual for information on solving possible interoperability problems.

2.4.5.2 Configuring a Secure Gateway

Before configuring IPsec on a router or a gateway, make sure that the system is configured as an IP router. See the *Network Administration: Connections* manual for information on configuring the system as an IP router.

To configure IPsec on a router or gateway, follow these steps:

1. From the SysMan Menu, select Networking→Additional Network Services→Set up IP Security (IPsec) to display the IPsec main window.

   Alternatively, enter the following command on the command line:

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

   If you are configuring IPsec for the first time, an informational dialog box is displayed that tells you to define secure connections before enabling IPsec. If you enable IPsec without defining secure connections, all packets into and out of the system are discarded; no traffic will flow. Select OK.

   The IPsec main window displays configured secure connections and configured public-key certificates.
2. Select Enable IP Security (IPsec) at the top of the window.
3. Select Add. The Add/Modify a Secure Connection dialog box is displayed.
4. Enter a connection name.
5. Select Add to add a remote IP address selectors. The Add/Modify Selector dialog box is displayed. Do the following:

   a. Select a selector type.

   b. Do one of the following:

      - If you are communicating with a single host, enter the IP address.
      - If you are communicating with a secure gateway, enter the subnet address.
      - If you are communicating with a range of addresses, enter the first address.

   c. For an IP subnet, enter the size of the subnet mask.

   d. For a range of addresses, enter the last address.

   e. Select an upper layer protocol to match. By default, all protocols are selected.

   f. Enter a port number to match, if you want to restrict the selector to a specific port number. By default, all port number are selected.

g. Select OK to accept the data and close the Add/Modify Selector dialog box. If you are finished selecting remote and local addresses, go to step 7.

6. Select Add to add a local IP address selector. Go to step 5a.
7. Select an action to apply to the packets matching the selectors. The default is to apply IPsec protection.
8. Select Next to accept the data and close the Add/Modify a Secure Connection dialog box. The Add/Modify Connection: IPsec Proposal dialog box is displayed. Do the following:

   a. Select an IPsec proposal from the proposal list.

   b. If you are communicating with a secure gateway or a host, specify the IP address of the remote system and your system's IP address (local).

   c. Specify if you will use IKE to obtain keys or use manual configuration. Select Next to accept the data and close the IPsec Proposal dialog box.

      If you selected manual configuration and have created a custom proposal list with only one proposal, the Add/Modify Connection: Manual Keys dialog box displays. Go to step 9. If you selected the IKE protocol, the Add/Modify Connection: IKE Proposal dialog box displays. Go to step 11.

9. Select Add to add a manual key and display the Manual Keys: Add/Modify IPsec Key dialog box. Do the following:

   a. Enter the key name.

   b. Enter the Security Parameter Index (SPI).

   c. Enter keys for the algorithms that are required by the proposals you chose. Select OK to accept the data and close the Manual Keys: Add/Modify IPsec Key dialog box.

10. Select whether you want to apply the key(s) to inbound packets, outbound packets, or both. If you want to specify additional keys, go to step 9. If you are finished specifying manual keys, select Finish. Go to step 20.
11. Select an IKE proposal from the proposal list.
12. Select Add to restrict access to the connection and display the Add/Modify Remote IDs dialog box. Do the following:

    a. Select a remote identity type.

b. Enter an identity string, usually your IP address, domain name, or email address.

c. Select OK to accept the data and close the Add/Modify Remote IDs dialog box.

Note

A remote identity (ID) is one that is allowed to use this connection. Identities are values that are either specified in a certificate by the Subject Alternate Name or that you enter when specifying a preshared key. This step is optional. However, if you do not specify a remote identity or identities, other systems might have unauthorized access to your system.

13. If you want to specify additional remote identities, go to step 12. If you are finished specifying remote identities, select Next to accept the data, close the Add/Modify Connection: IKE Proposal dialog box, and display the Add/Modify Connection: IKE Authentication dialog box.
14. Select whether you want to authenticate IKE exchanges with a public-key certificate or a preshared key.
15. If you selected public-key certificate, select Add to add an IKE certificate. The Add/Modify Certificates dialog box is displayed. Do the following:

    a. Enter a certificate name, select a certificate encoding method, and enter the local path to the certificate file.

    b. If the certificate authenticates your system, select the encoding method and enter the local path to the private key file.

    c. If the certificate is trusted to sign other certificates, select CA Certificate. Otherwise, go to step f.

    d. If a Certificate Revocation List (CRL) is not available, select No Certificate Revocation List (CRL) Available. Go to step f.

    e. Select an encoding method for the CRL and enter a local path to the CRL file.

    f. Select OK to accept the data and close the Add/Modify Certificates dialog box.

16. Select a certificate for the IKE exchange. Go to step 19.
17. If you selected preshared key, select Add an IKE preshared key. The Add/Modify IKE Keys dialog box is displayed. Do the following:

    a. Enter a key name and key value.

b. Select a local identity type.

c. Enter an identity string, usually your IP address or domain name.

d. Select OK to accept the data and close the Add/Modify IKE Keys dialog box.

18. Select a preshared key for the IKE exchange.

19. Select Next to close the Add/Modify Connection: IKE Authentication dialog box and display the Add/Modify Connection: Optional IKE Parameters dialog box. Do the following:

    a. Select any optional parameters.

    b. Select an IKE group number for initial Diffie-Hellman exchanges, if it is different from the IKE proposals.

    c. If using Perfect Forward Secrecy (PFS), select a group number future for Diffie-Hellman exchanges.

    d. Select a default lifetime if the proposal does not specify a lifetime.

    e. Select Finish to accept the data and close the Add/Modify Connection: Optional IKE Parameters dialog box.

20. An informational dialog box is displayed that tells you the connection has been created. Select OK to close this dialog box.

21. If you need to specify additional public-key certificates, select Add in the Public-Key Certificates field to display an Add/Modify Certificates dialog box into which you can enter information for the certificate. Do the following:

    a. Enter the certificate name, select a certificate encoding method, and enter a local path to the certificate file.

    b. If the certificate authenticates your system, select a private key encoding method and enter a local path to the private key file.

    c. If the certificate is trusted to sign other certificates, select CA Certificate. Otherwise, go to step f.

    d. If a Certificate Revocation List (CRL) is not available, select No Certificate Revocation List (CRL) Available. Go to step f.

    e. Select an encoding method for the CRL and enter a local path to the CRL file.

    f. Select OK to accept the data and close the Add/Modify Certificates dialog box.

22. Select OK in the IPsec main window to save the configuration information. Whether or not IPsec is already running on your system,

the Restart IPsec? dialog box is displayed. If you want to start or restart IPsec, select OK; otherwise, select No. If you select No, you can reboot the system to start or restart IPsec, or start or reload the `ipsecd` daemon (see the Network Administration: Connections manual).

See the *Network Administration: Connections* manual for information on solving possible interoperability problems.

2.4.6 Adding Callout Functions for IP Processing

The `fr_checkp` global variable is a callout hook in the kernel IP processing code. You can use this hook to call out to a customized routine to filter or verify IP packets.

To add a callout in the IP input and output processing, create a module that performs an assignment of `fr_checkp` during the initialization or configuration of the custom filter module, as follows:

```
(*fr_checkp) (struct ip *ip, int hlen, struct ifnet *rcvif, int
direction, struct **mbuf bufp)
```

Where:

ip	Points to the IP header.
hlen	Is the length of the header.
rcvif	Is a pointer to the receiving or sending interface.
direction	0 for input; 1 for output.
bufp	Is a pointer to the `mbuf` message chain.

If the routine returns a zero, IP processing continues using the `mbuf` pointer returned in the *bufp* field. If a nonzero value is returned or if the `mbuf` pointer is zero, IP processing is terminated.

If the callout function returns a nonzero value, the callout routine must free the `mbuf` chain using `m_freem`.

The following example shows how to create a module, `custom_filter`, which filters out a packet if it matches the selected type of service (TOS) field of the IP header:

```
#include "sys/errno.h"
#include "net/if.h"
#include "netinet/ip.h"
#include "sys/mbuf.h"
#include "sys/sysconfig.h"
```

```
char custom_filter_tos = 255;
static int debug=0;
char custom_filter_version[] = "custom_filter: V1.00";

cfg_subsys_attr_t packetfilter_attributes[]  = {
        /*
         * name of the table
         */
        {"version", CFG_ATTR_STRTYPE,
         CFG_OP_QUERY,
         (caddr_t) custom_filter_version, 2, 100, 0},
        /*
         * debug state
         */
        {"debug", CFG_ATTR_ULONGTYPE,
         CFG_OP_CONFIGURE | CFG_OP_QUERY | CFG_OP_RECONFIGURE,
         (caddr_t) &debug, 0, ULONG_MAX, 0},

        /*
         * Tos to filter on
         */
        {"tos", CFG_ATTR_UCHARTYPE,
         CFG_OP_QUERY | CFG_OP_CONFIGURE,
         (caddr_t) &custom_filter_tos, 0, 255, 0},

        {"", 0, 0, 0, 0, 0, 0}  /* must be the last element */ };

int
custom_filter(struct ip *ip, int hlen, struct ifnet *rcvif,
       int direction, struct mbuf **bufp)
{
        if( ip->ip_tos == custom_filter_tos ){
                mfreem(bufp);
                return(1);
        }
        return(0);
}

custom_filter_configure(
        cfg_op_t                op,
        caddr_t                 indata,
        ulong                   indata_size,
        caddr_t                 outdata,
        ulong                   outdata_size)
{
 extern int (*fr_checkp) (struct ip *ip, int hlen, struct ifnet *rcvif,
     int direction, struct mbuf **mbuf);
        switch (op) {
        case CFG_OP_CONFIGURE:

                fr_checkp=custom_filter;
                break;
        case CFG_OP_UNCONFIGURE:
                fr_checkp=NULL;
        }
        if( debug > 1 )
                printf("custom_filter_configure: returning ESUCCESS\n");
        return ESUCCESS;
}
```

3

New and Changed Features for Tru64 UNIX Version 5.1B-1

This chapter describes new features that are available with the Version 5.1B-1 release of the operating system. It also lists new hardware that is supported and provides information about retiring products.

3.1 Support for Tuning Big Pages Attributes of Binary Files

This release provides support for tuning binary files to have different big page behavior than system defaults. These settings can override the system defaults for specific types of memory (anonymous, program text, SysV shared, SysV shared segmented, and stack).

Each of the specific type settings has system-wide tunables, expressed as a threshold in Kbytes. The default is 64 Kbytes, the size of the smallest big page. The per-binary tunables are also expressed as a threshold in Kbytes.

An additional tunable directs big pages to distribute memory across RADs as a priority over getting the largest page size possible.

For information about using this feature, see `chatr`(1) and `sys_attrs_vm`(5) and the *System Configuration and Tuning* guide.

3.2 Support for the Name Services Switch

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 the `/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 `nssetup`(8), `nsswitch`(4), and `nss2svc`(8).

3.3 New Security Feature

Patch 1414.00 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, `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 optionally to generate a message when such a request occurs.

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

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

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

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

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

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

Note

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

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

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

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

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

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

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

See chatr(1) for more information.

3.4 Packetfilter Enhancements

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

- pfilt_loopback=[0|1] Setting this flag to 0 prevents a loop back of any packetfilter-written multicast or broadcast packet. When set to 1, the packetfilter loops back both broadcast and multicast packets. This is the default.
- pfilt_physaddr=[0|1] Setting this flag to 0 allows the application to fill in the source Ethernet address for packetfilter-written packets. If the source address is all 0s, the address is set to the proper hardware address by the packetfilter code as a safety precaution. When set to the packetfilter sets the Ethernet source address in the outgoing packet. This is the default.

You can set these variables at boot time as follows:

```
# echo " patch pfilt_loopback=0 " | dbx -k /vmunix
# echo " patch pfilt_hysaddr=0 " | dbx -k /vmunix
```

Restriction

The pfilt_loopback and pfilt_physaddr tunable variables are only accessible in the kernel using dbx. In a future patch kit, these tunable variables will be implemented as kernel subsystem configurable attributes, accessible through the sysconfig command.

3.5 New Hardware Support

The following new hardware support is provided with this patch kit.

3.5.1 Support for 64–Processor AlphaServer GS1280 Systems

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

3.5.2 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

3.5.3 HP StorageWorks FCA2384

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

3.6 New and Updated Associated Products

A number of software products provided on the Associated Products CD–ROMs have been updated for this release. The updated products are listed in the following sections.

For more information on the CD–ROM contents, see the *HP Tru64 UNIX Version 5.1B-1 CD-ROMs* card contained in the media kit.

3.6.1 Advanced Printing Software Version 1.2A

Advanced Printing Software has been updated to support the following printers:

Genicom mL210 PS
Genicom mL280 PS
Genicom LN21 PS
Genicom LN28 PS
Genicom cL160 PS
HP LaserJet 2300 Series Printers PS
HP LaserJet 4200 Series Printers PS
HP LaserJet 4300 Series Printers PS
HP LaserJet 5100 Series Printers PS
HP LaserJet 2500 Series Color Printers PS
HP LaserJet 4550 Series Color Printers PS
HP LaserJet 4600 Series Color Printers PS
HP LaserJet 5500 Series Color Printers PS

The list of supported printers can be found at the following URL:
http://h30097.www3.hp.com/printing/printers.html

3.6.2 Advanced Server for Tru64 UNIX

Advanced Server for Tru64 UNIX (ASU) has been updated to Version 5.1B ECO1. This update includes support for the following new features:

- The ASU server supports systems configured with LAG (Link aggregation) network support. You can configure the ASU server to listen on NetBIOS over TCP/IP and NetBEUI over the LAG interface.

- The ability to enable and disable event logging by the Event Manager (EVM).
- When negotiating a protocol with a remote server, the ASU server now sends the list of dialects that it supports in one SMB packet. This improves ASU server performance and keeps the lmx.dmn process from hanging when the remote server is Windows® NT®.

For a complete description of the changes made to ASU, read the release notes available at the following URL:
http://h30097.www3.hp.com/docs/asdu/HTML/asdu.html

3.6.3 Application Transition Tools

The following Tru64 UNIX to HP-UX application transition tools have been added to the Associated Products CD–ROM:

- Tru64 UNIX to HP-UX Software Transition Kit Version 2.0 – This kit includes file scanning utilities, developer's documentation, and porting documentation to help resolve compatibility issues between Tru64 UNIX and HP-UX. The file scanning utilities use a clear methodology for code analysis, providing sound advice for each Application Programming Interface (API) encountered in scanned Tru64 UNIX C and C++ source code files.
- appscan Version 2.0 – This utility enables you to list all of the dependencies (shared libraries and symbols) of a dynamic executable file. The utility also generates an associated disposition code for each of the listed APIs.
- hpuxman v1.0 – This command allows you to display select HP-UX 11i v1.6 reference pages on systems running the Tru64 UNIX operating system.

3.6.4 Compaq COBOL RTL

The Compaq COBOL RTL has been updated from Version 2.7 to Version 2.8-670.

3.6.5 OpenLDAP Directory Server

This is an update of the OpenLDAP Directory Server from Version 2.0.23 to Version 2.0.27. This update consists mainly of bug fixes and includes new versions of all of the OpenSource components.

3.6.6 OpenLDAP Utilities

The LDAP Client Utilities have been updated to Version 1.1. This update includes support for an additional configuration parameter, `nisnetgrpbranch`, added for LDAP netgroups support.

3.6.7 Mozilla Version 1.4 Application Suite for Tru64 UNIX

The Mozilla Application Suite is the next generation Web, mail, and news application successor to the popular Netscape Communicator Web client. Mozilla is an open source Web application created by the Mozilla Foundation. It is designed for standards compliance, performance, and portability.

The Mozilla 1.4 Application Suite also includes many new innovative features for search, privacy, and content management of your Internet information. Built upon the Netscape Gecko browser engine, the Navigator component is now comprehensive, modular, and fully standards compliant — supporting DOM, RDF, XML, CSS, and HTML 4 document formats.

3.6.8 Java

Java has been updated from Version 1.3.1 to Version 1.4.1–2.

3.6.9 Secure Web Server

Secure Web Server has been updated to Version 6.1. This update contains a new subset based on Apache 2 in addition to the older Version 1.3 code base. The update includes new versions of all of the OpenSource components.

3.6.10 Legato NetWorker

Legato NetWorker for HP Tru65 UNIX has been updated to Version 7.0. This update includes a number of new and advanced features.

3.6.11 WEBES

The Web Based Enterprise Service Suite has been updated to Version 4.2.

3.6.12 Unicensus

The Unicensus Revision and Configuration Management (RCM) application has been updated from Version 4.5.2 to Version 4.5.4.

3.6.13 Visual Threads

Visual Threads Version 2.3 offers the following new features:

- Enhancements to profiling and dynamic resizing of the Event Details window
- Fixes for the Summary and Print windows
- Fixes for deadlock and inconsistent Order detection

3.7 Sources for Open Source Components

In this release, a new CD–ROM has been added to the Tru64 UNIX media kit. This CD–ROM contains sources for Open Source Software provided with the Tru64 UNIX operating system.

The *Sources for Open Source Components* CD–ROM contains sources for the following products:

- Secure Web Server:

  Secure Web Server Administration Utility
  Secure Web Server Documentation
  Tomcat Java Servlet and JSP Engine
  Secure Web Server 1.3 powered by Apache 1.3
  Secure Web Server 2.0 powered by Apache 2.0

- LDAP Components

  OpenLDAP Directory Server
  LDAP Client Utilities

- Mozilla 1.4 for Tru64 UNIX

  Mozilla 1.4 Application Suite
  Mozilla 1.4 Runtime Support

3.8 Retirement Notices

This section provides information on features that have been retired from the operating system.

3.8.1 Aurema ARMTech Products Retirement

Aurema ARMTech products, ShareExpress, ShareExtra, and ShareEnterprise will be removed from the Tru64 UNIX operating system distribution before December 2003. HP will continue to be support ShareExpress through June 2004.

Aurema announced end of sales of the ShareExtra and ShareEnterprise products for Tru64 UNIX in June 2003. Aurema continues to directly support current customers.

3.8.2 DEC Ada Retirement

Section 2.2.8 of the Version 5.1B *Release Notes* incorrectly states that DEC Ada (UPI - 0HM) and DEC Ada PDO (UPI - 0VS) will be retired in a future release of the operating system.

DEC Ada was retired in March 2000.

4

Software Notes for Tru64 UNIX Version 5.1B-1

This chapter contains notes about issues and known problems with the operating system and, whenever possible, provides solutions or workarounds to those problems. The chapter also describes any changes or corrections to the documentation since the last release of the operating system.

4.1 Software Notes and Restrictions

The following topics describe restrictions and known problems and workarounds to the operating system software.

4.1.1 Problem with the find Command

The `find` command may fail when traversing large directory structures. Messages similar to the following will be displayed when this problem is encountered:

```
$ find . -name abc
find: Cannot open file ./aa5142
find: Cannot open file ./aa5143
find: Cannot open file ./aa5144
:Too many open files
```

HP has corrected this problem and has provided the fix in an Early Release Patch Kit (ERP). The ERP kit name is T64KIT0020545-V51BB24-E-20031104 and can be accessed at the following URL:
http://www.itrc.hp.com/service/patch/patchDetail.do?patchid=T64KIT0020545-V51BB24-E-20031104

4.1.2 Problem Displaying Apache Documentation

The following text is displayed in the Web browser window when the Apache Documentation link is clicked on the Secure Web Server 2.0 initial top level Web page:

```
URI: index.html.de .... charset=ISO-8859-1
   URI: index.html.en ....  charset=ISO-8859-1
   URI: index.html.fr ....  charset=ISO-8859-1
   URI: index.html.ja.....  charset=ISO-2022-JP
   URI: index.html.ko.euc-kr .. charset=EUC-KR
```

To resolve this problem, execute the following script as root from a terminal window or the console prompt. This script adds the .var filename extension to all of the files under the/usr/opt/hpapache2/manual directory tree with an .html filename extension and that contain the string 'URI: ' as the first characters in a line of the file:

```
#!/sbin/sh
   # convert the multiviews files to .var files
   cd /usr/opt/hpapache2/manual
   find . -name '*.html' | xargs grep -lE '^URI: ' | while read file
   do
      mv $file ${file}.var
   done
```

After running the script, click on the Apache Documentation link again in a Web browser and the Apache HTTP Server Version 2.0 Documentation Web page will appear.

4.1.3 Potential NFS Duplicate Request Cache Scalability Limitation 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 undistinguishable from the server going up and down, except that the messages are displayed with such frequency that the server system/member cannot have gone down and then come back up in that short of 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 retransmisions. A retransmission rate higher than 2 percent indicates a potential problem.

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

```
% nfsstat -c

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

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

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

- Ensure that there are short periods of idle time on the I/O subsytem 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 a 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 whether 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
```

4.1.4 Tuning the NFS Server Duplicate Request Cache

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

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

Patch 1062.00 provides a tuning variable to control the size of the NFS server's duplicate request cache: `nfs_dupcache_size`. This variable controls the absolute size of the NFS server duplicate request cache. This is measured in the number of elements that are allocated at NFS server initialization.

If the size of the duplicate cache needs to be modified, change `nfs_dupcache_size`. Set the new value for `nfs_dupcache_size` 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.

4.1.5 Problems Uninstalling the Patch Kit

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

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

To uninstall the patch kit, do the following:

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

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

4.1.6 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 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.

4.1.7 Possible Error Seen with Patch 1276.00

After installing Patch 1276.00, the following problems have been known to occur:

- 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/
```

4.1.8 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 -tlr /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 the cluster software misinterpreting the existence of some patches and will be corrected in a future patch kit.

If the `ls -tlr PAT.sts` 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.

4.1.9 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
```

Do not use his kernel setting as a default. Enable it only 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 `xargs`(1).

4.1.10 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. The subset where the `bootpd` ships is optional, so OSFOBSOLETE540 must be installed.

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

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

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

```
 # Example bootptab file for BOOTP support

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

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

In this example, the directory `/install/firmware` was created on the `bootp` server. This directory must contain the firmware of the systems to be

updated. The file names must match the entry on the fwupgrade command line. The firmware files must have read permissions, that is, 444.

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

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

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

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

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

4.1.11 Broken Links Reported During Baselining

When performing a baseline analysis with the dupatch utility, you can disregard the following message during Phase 4:

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

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

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

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

    Press RETURN to proceed...
```

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.

4.1.12 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 have been supported since Patch Kit 1 for Tru64 UNIX Version 5.1B was released.

4.1.12.1 Memory Restriction on 64–Processor GS1280 Systems

In AlphaServer GS1280 systems with 64 processors and various memory sizes, you must configure the largest amount of memory in RAD0.

4.1.12.2 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 systems do not allow any CPUs to be taken off line.

4.1.12.3 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. After a CPU is marked as spare, the `hwmgr` command and Manage CPUs SUITLET prevent you from putting them on line until you use the `ccod -l` or `ccod -p` command either to loan or to purchase the CPU.

To work around this problem, 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 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.

4.1.12.4 Hardware SCSI Errors

SCSI errors experienced by the Adaptec controller that require SCSI bus resets can cause PCI bus faults. These faults will be seen as a “Machine Check System Uncorrectable” panic. This will require the system to be

booted after the machine check. A fix for this problem will be included in a future release.

4.1.12.5 Compact Disk Drive Errors Logged

The TEAC CDR-W 416E drive that is shipped with the AlphaServer system will log errors on reboot if the CD-ROM media is not present. These messages are only informational.

4.1.12.6 Presence of Third-Party Devices May Cause System Panic

The ATM 3X-DAPBA-FA/UA driver may experience a panic on shutdown if third-party devices are installed.

4.1.12.7 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.

4.1.12.8 Incorrect Free Page Counts Reported

The `vmstat -S` command reports incorrect free page counts on a sparsely configured system. A sparsely configured system has gaps in the numbering of CPUs; for example, 0, 1, 8, 9, 10, 11.

4.1.12.9 Increasing PCI Box Support on GS1280 M32 to 16

The maximum number of standard I/O hoses (IDE buses) allowed for an initial install of Tru64 UNIX Version 5.1B from the CD-ROM is eight. Any I/O drawer connected to a hose that has more than eight must be disconnected or powered down during a fresh install from the Tru64 UNIX CD-ROM.

4.2 Documentation Notes

4.2.1 AdvFS Administration Manual Correction — Extend an AdvFS File System When Increasing the Size of the Underlying Volume

In Section 2.3.4.3 of the *AdvFS Administration* manual — "Increasing Storage in Domains by Extending an Existing Volume" — a step is missing from the procedure when the underlying storage volume is a hardware RAID device. You must modify the volume's disk label information to reflect the new, increased size of the partition supporting the domain, and then apply the updated disk label to the volume before extending the file system.

The complete process to extend a domain by increasing the size of an underlying hardware RAID volume includes the following steps:

1. Using HSG80 commands, extend the hardware RAID volume.

   This might involve adding another stripeset to an existing stripeset, or creating a concatset from the original hardware RAID volume and adding another volume to it.

   For example, assume the AdvFS domain uses disk `dsk25c`, which is a single hardware RAID volume. To extend the capacity of the disk, create a concatset from it and another single hardware RAID volume, as shown in the following example:

```
HSG80> show disks [1]
Name          Type                        Port Targ  Lun        Used by
------------------------------------------------------------------------------

DISK10000     disk                           1    0    0        DELI-5.1A
DISK10100     disk                           1    1    0        SPARESET
DISK10200     disk                           1    2    0        D8          [2]
DISK20000     disk                           2    0    0        DELI-5.1A
DISK20100     disk                           2    1    0        SPARESET
DISK20200     disk                           2    2    0                    [3]
DISK30000     disk                           3    0    0        GALLO-M1
DISK30200     disk                           3    2    0
DISK40000     disk                           4    0    0        GALLO-M1
DISK40200     disk                           4    2    0
DISK50000     disk                           5    0    0        GALLO-M2
DISK50200     disk                           5    2    0
DISK60000     disk                           6    0    0        GALLO-M2
DISK60200     disk                           6    2    0
HSG80> add concatsets C1 DISK10200 [4]
HSG80> set C1 add=DISK20200 [5]
HSG80> show C1 [6]
Name          Storageset                    Uses             Used by
------------------------------------------------------------------------------

C1            concatset                     DISK10200        D8
                                            DISK20200

        State:
          NORMAL
          DISK10200 (member  0) is NORMAL
          DISK20200 (member  1) is NORMAL
        Size:             71112778 blocks [7]
```

   The following list explains each step:

[1] Find unused disks — those with an empty (blank) `Used by` field.

[2] `DISK10200` (also called `D8`) is used by the AdvFS domain that will be extended. This storage volume is recognized as `dsk25` on Tru64 UNIX.

[3] `DISK20200` is unused.

[4] Create a concatset called `C1` from `DISK10200`.

[5] Add `DISK20200` to concatset `C1` to create a larger disk.

[6] Display the size of the concatset `C1`.

7 Note the size, because you will need to modify the disk label for `dsk25` on Tru64 UNIX to match it.

2. Return to the Tru64 UNIX prompt.
3. Save a copy of the old disk label information for the volume:

   # **disklabel -r dsk*N* > /tmp/*label***

   For example:

   # **disklabel -r dsk25 > /tmp/dsk25MOD**

4. Edit the saved label and increase the size of the partition used by AdvFS:

   # **vi /tmp/*label***

   For example:

   # **vi /tmp/dsk25MOD**

5. Write the edited disk label back to the hardware RAID volume:

   # **disklabel -R dsk*N* /tmp/*label***

   For example:

   # **disklabel -R dsk25 /tmp/dsk25MOD**

6. Optionally, display the size of the domain before extending it:

   # **showfdmn *domain***

   For example, for a domain called `clinical_trials`, enter:

   # **showfdmn clinical_trials**

```
             Id              Date Created  LogPgs  Version  Domain Name
3e8ca76d.040d38c8  Thu Apr  3 16:28:13 2003     512        4  clinical_trials

  Vol   512-Blks        Free  % Used  Cmode  Rblks  Wblks  Vol Name
   1L   35556384    35547280      0%     on    256    256  /dev/disk/dsk25c
```

7. Extend the AdvFS file system:

   # **mount -u -o extend */file_system***

   For example, if the AdvFS domain that uses `dsk25` is mounted on `/test_data`, enter:

   # **mount -u -o extend /test_data**

8. Optionally, verify that the domain now shows the larger size:

   # **showfdmn *domain***

   For example:

   # **showfdmn clinical_trials**

```
              Id              Date Created  LogPgs  Version  Domain Name
3e8ca76d.040d38c8  Thu Apr  3 16:28:13 2003     512        4  clinical_trials

  Vol   512-Blks        Free  % Used  Cmode  Rblks  Wblks  Vol Name
```

```
1L   71112768    71103120      0%     on    256    256  /dev/disk/dsk25c
```

4.2.2 Installation Guide Contains Incorrect Java Version

The *Installation Guide* states that Version 1.3.1-1 of Java is provided with Version 5.1B of the operating system. This is incorrect. Java Version 1.3.1-2 is provided with the Version 5.1B release of the operating system.

Version 5.1B-1 of the operating system includes Java Version 1.4.1-2.

4.2.3 System Configuration and Tuning Guide Corrections

Section 4.2 of the *System Configuration and Tuning* guide provides an example demonstrating how to enable access to the system's real time clock. This example is incorrect. The correct command is:

```
# mknod /dev/timedev c 15 0
```

Section 4.4.8.4 states: "The `max_async_req` attribute specifies the maximum number of sessions within any given RDG context table. The recommended value is at least the number of Oracle® processes plus two." This is incorrect. The `max_sessions` attribute specifies the maximum number of sessions within any given RDG context table.

Section 4.4.8.5 states: "The `max_async_req` attribute specifies the maximum number of pages automatically wired in memory for message packets." This is incorrect. The `rdg_max_auto_msg_wires` attribute specifies the maximum number of pages automatically wired in memory for message packets. We recommend setting this attribute to 0.

Section 6.2.2.2 states that, if you increase the value of the `max_proc_per_user` attribute, you increase the amount of wired memory. This statement is false. Increasing this attribute value does not increase the amount of wired memory.

4.2.4 ypset(8) Correction

The `ypset`(8) reference page indicates that both V1 and V2 are allowed as options. This is incorrect. V2 is not a supported option.

4.2.5 aio_return(3) Correction

The following information should be appended to the RETURN VALUES section of `aio_return`(3):

On an unsuccessful call, the value of -1 is returned and `errno` is set to indicate the error. If the operation did not complete, but it terminated normally (because, for example, the call was purposely interrupted by the `aio_cancel` function), `errno` is set to 0.

4.2.6 disklabel(8) Correction

The following definition of the output from the disklabel command is missing from recent versions of disklabel(8):

An asterisk (*) is sometimes shown in the output from the disklabel command, under the column headed cylinders grouped for a partition (cpg).

This asterisk indicates that the start or the end of a cylinder does not fall exactly on a block boundary.

4.2.7 dxshutdown(8) Correction

Since the release of Tru64 UNIX Version 5.0, the command /usr/bin/X11/dxshutdown is a wrapper shell script that runs the SysMan shutdown program. Prior to Version 5.0, dxshutdown was an X motif application. The X motif version of dxshutdown is shipped in an obsolete subset. The new dxshutdown shell script can run the old version when it is installed as /usr/bin/X11/dxshutdown_old. Use the following command:

```
# /usr/bin/X11/dxshutdown -old
```

The current OPTIONS section of dxshutdown(8) is no longer applicable because suitlets use Tk and Tk uses X, not Xt. The only useful argument is -focus *hostname* when running on a cluster.

The SysMan application is no longer called Shutdown Manager. If invoked from the SysMan menu, the leaf is labelled "Shutdown the system" and the application is labelled "Shutdown targeted on *hostname*".

In the EXAMPLES section, the /usr/dt/appconfig/help/C/DXshutdown.sdl help file is no longer used. In the FILES section, /usr/dt/appconfig/help/C/Dxshutdown.sdl and $HOME/Dxshutdown are no longer used.

4.2.8 emx(7) Correction

The emx(7) reference page provides an example of how to turn off I/O limiting by using the following run-time configuration command:

```
# /sbin/sysconfig -r io NPort_Max_IOs = 0xFFFFFFFF
```

This example command should read as follows:

```
# /sbin/sysconfig -r emx NPort_Max_IOs=0xFFFFFFFF
```

4.2.9 dd(1) Correction

Note 1 in dd(1) provides an example of zeroing a disk label. The syntax of the disklabel command used in that example is incorrect and should read as follows:

```
# disklabel -r /dev/rdisk/dsk1a
# disklabel -z  /dev/rdisk/dsk1a
# disklabel: Disk /dev/rdisk/dsk1a is unlabeled
```

4.2.10 ksh(1) correction

In ksh(1), the following statements are incorrect:

&& Causes the list following it to be executed only if the preceding pipeline returns a 0 (zero) exit value.
|| Causes the list following it to be executed only if the preceding pipeline returns a nonzero exit value.

The correct statements are:

&& Causes the list following it to be executed only if the preceding pipeline returns a nonzero exit value.
|| Causes the list following it to be executed only if the preceding pipeline returns a 0 (zero) exit value.

Index