

Tru64 UNIX

Release Notes for Version 5.1B-1

Part Number: AA-RVGEA-TE

November 2003

Product Version: Tru64 UNIX Version 5.1B-1

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 Tru64 UNIX Version 5.1B *Technical Update* 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/pub_page/update_list.html

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

Microsoft®, Windows®, and Windows NT® are U.S. registered trademarks of Microsoft Corporation. Intel® is a registered trademark of Intel Corporation or its subsidiaries in the United States and other countries. Java™ is a U.S. registered trademark of Sun Microsystems, Inc. Oracle® is a registered U.S. trademark of Oracle Corporation, Redwood City, California. Motif®, OSF/1®, UNIX®, and X/Open® are registered trademarks and The Open Group™ is a trademark of The Open Group in the U.S. and other countries. All other product names mentioned herein may be trademarks of their respective owners.

Proprietary computer software. Valid license from HP required for possession, use, or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Neither HP nor any of its subsidiaries shall be liable for technical or editorial errors or omissions contained herein. The information is provided "as is" without warranty of any kind and is subject to change without notice. The warranties for HP products are set forth in the express limited warranty statements accompanying such products. Nothing herein should be construed as constituting an additional warranty.

Contents

About This Manual

1 New and Changed Features

1.1	Support for Tuning Big Pages Attributes of Binary Files	1-1
1.2	Support for the Name Services Switch	1-1
1.3	New Security Feature	1-1
1.4	Packetfilter Enhancements	1-4
1.5	New Hardware Support	1-4
1.5.1	Support for 64-Processor AlphaServer GS1280 Systems ..	1-4
1.5.2	Support for AlphaServer and AlphaStation DS15 Systems	1-4
1.5.3	HP StorageWorks FCA2384	1-5
1.6	New and Updated Associated Products	1-5
1.6.1	Advanced Printing Software Version 1.2A	1-5
1.6.2	Advanced Server for Tru64 UNIX	1-5
1.6.3	Application Transition Tools	1-6
1.6.4	Compaq COBOL RTL	1-6
1.6.5	OpenLDAP Directory Server	1-6
1.6.6	OpenLDAP Utilities	1-7
1.6.7	Mozilla Version 1.4 Application Suite for Tru64 UNIX	1-7
1.6.8	Java	1-7
1.6.9	Secure Web Server	1-7
1.6.10	Legato NetWorker	1-7
1.6.11	WEBES	1-7
1.6.12	Unicensus	1-7
1.6.13	Visual Threads	1-7
1.7	Sources for Open Source Components	1-8
1.8	Retirement Notices	1-8
1.8.1	Aurema ARMTech Products Retirement	1-8
1.8.2	DEC Ada Retirement	1-9

2 Software Notes

2.1	Software Notes and Restrictions	2-1
2.1.1	Problem with the find Command	2-1
2.1.2	Problem Displaying Apache Documentation	2-1

2.1.3	Potential NFS Duplicate Request Cache Scalability Limitation on Clustered NFS Servers	2-2
2.1.4	Tuning the NFS Server Duplicate Request Cache	2-4
2.1.5	Problems Uninstalling the Patch Kit	2-5
2.1.6	Performance of hwmgr Commands on Large System Configurations	2-5
2.1.7	Possible Error Seen with Patch 1276.00	2-6
2.1.8	Error on Cluster Creation	2-6
2.1.9	Possible Problem When Processing Many Command Parameters	2-7
2.1.10	Loading Firmware from a BOOTP Server	2-7
2.1.11	Broken Links Reported During Baselineing	2-8
2.1.12	General and Problem Information for AlphaServer ES47, ES80, and GS1280 Systems	2-8
2.1.12.1	Memory Restriction on 64-Processor GS1280 Systems	2-8
2.1.12.2	CPU Offline Restrictions	2-8
2.1.12.3	Problem with Capacity-on-Demand Process	2-9
2.1.12.4	Hardware SCSI Errors	2-9
2.1.12.5	Compact Disk Drive Errors Logged	2-9
2.1.12.6	Presence of Third-Party Devices May Cause System Panic	2-10
2.1.12.7	Repeated Reboots May Cause Panic	2-10
2.1.12.8	Incorrect Free Page Counts Reported	2-10
2.1.12.9	Increasing PCI Box Support on GS1280 M32 to 16	2-10
2.2	Documentation Notes	2-10
2.2.1	AdvFS Administration Manual Correction — Extend an AdvFS File System When Increasing the Size of the Underlying Volume	2-10
2.2.2	Installation Guide Contains Incorrect Java Version	2-12
2.2.3	System Configuration and Tuning Guide Corrections	2-13
2.2.4	ypset(8) Correction	2-13
2.2.5	aio_return(3) Correction	2-13
2.2.6	disklabel(8) Correction	2-13
2.2.7	dxshutdown(8) Correction	2-14
2.2.8	emx(7) Correction	2-14
2.2.9	dd(1) Correction	2-14
2.2.10	ksh(1) correction	2-15

Index

About This Manual

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

This manual also describes significant new and changed features in this version 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 this version of the operating system software
- Chapter 2* Contains information on restrictions to the software and documentation

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 or hardcopy 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 Tru64 UNIX Version 5.1B *Technical Update* 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/pub_page/update_list.html

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:

<code>%</code>	
<code>\$</code>	A percent sign represents the C shell system prompt. A dollar sign represents the system prompt for the Bourne, Korn, and POSIX shells.
<code>#</code>	A number sign represents the superuser prompt.
<code>% cat</code>	Boldface type in interactive examples indicates typed user input.
<code><i>file</i></code>	Italic (slanted) type indicates variable values, placeholders, and function argument names.
<code>[]</code> <code>{ }</code>	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.
<code>...</code>	In syntax definitions, a horizontal ellipsis indicates that the preceding item can be repeated one or more times.
<code>cat(1)</code>	A cross-reference to a reference page includes the appropriate section number in parentheses. For example, <code>cat(1)</code> indicates that you can find information on the <code>cat</code> command in Section 1 of the reference pages.
<code>Ctrl/x</code>	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, <code>Ctrl/C</code>).

New and Changed Features

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

1.1 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 `chattr(1)` and `sys_attrs_vm(5)` and the *System Configuration and Tuning* guide.

1.2 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)`.

1.3 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 `chattr` utility, included in Patch 872.00 and described as follows:

```
$chattr +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 `chattr(1)` for more information.

1.4 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.

1.5 New Hardware Support

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

1.5.1 Support for 64-Processor AlphaServer GS1280 Systems

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

1.5.2 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

1.5.3 HP StorageWorks FCA2384

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

1.6 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.

1.6.1 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>

1.6.2 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>

1.6.3 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.

1.6.4 Compaq COBOL RTL

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

1.6.5 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.

1.6.6 OpenLDAP Utilities

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

1.6.7 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.

1.6.8 Java

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

1.6.9 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.

1.6.10 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.

1.6.11 WEBES

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

1.6.12 Unicensus

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

1.6.13 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

1.7 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

1.8 Retirement Notices

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

1.8.1 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.

1.8.2 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.

2

Software Notes

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.

2.1 Software Notes and Restrictions

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

2.1.1 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>

2.1.2 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.

2.1.3 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 retransmissions. A retransmission rate higher than 2 percent indicates a potential problem.

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

```
% nfsstat -c

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

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

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

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

- Use 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

2.1.4 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.

2.1.5 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.

2.1.6 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.

2.1.7 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/
```

2.1.8 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 -tldr /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 -tldr 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.

2.1.9 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 this 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)`.

2.1.10 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/bootpd
```

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.

2.1.11 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.

2.1.12 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.

2.1.12.1 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`.

2.1.12.2 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.

2.1.12.3 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 SUTLET 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 SUTLET 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 SUTLET is available from the SysMan Menu and SysMan Station.

2.1.12.4 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.

2.1.12.5 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.

2.1.12.6 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.

2.1.12.7 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.

2.1.12.8 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.

2.1.12.9 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.

2.2 Documentation Notes

2.2.1 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 dskN > /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 dskN /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

2.2.2 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.

2.2.3 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.

2.2.4 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.

2.2.5 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.

2.2.6 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 (`cpgr`).

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

2.2.7 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.

2.2.8 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
```

2.2.9 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
```


2.2.10 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

Numbers and Special Characters

64-Processor AlphaServer
systems, 1-4, 2-8

A

Advanced Printing Software
supported printers, 1-5

Apache documentation link, 2-1

ASU

new features, 1-5

Aurema ARMTech Products

retirement, 1-8

B

Big pages

tuning attributes in binary files, 1-1

BOOTP Server

firmware upgrade, 2-7

buffer overflow protection, 1-1

C

cache scalability limitation, 2-2

Cluster creation

error after deleting patches, 2-6

COBOL RTL

update, 1-6

D

DEC Ada

retirement, 1-9

dupatch

errors during baseline analysis, 2-8

E

EVM

support with ASU, 1-5

F

Firmware upgrade

using BOOTP server, 2-7

H

hwmgr

command performance on large
systems, 2-5

J

Java

privileged application failure, 1-2

update, 1-7

version number correction, 2-12

L

LDAP

as a source for netgroup data, 1-1

directory server update, 1-6

sources, 1-8

utilities update, 1-7

Legato NetWorker

update, 1-7

Link aggregation

support with ASU, 1-5

M

Memory

configuration restriction in GS1280 systems, 2-8
protecting against buffer overflow exploits, 1-2
tuning big page attributes, 1-1

Mozilla

application suite support, 1-7
sources, 1-8

N

Name Service Switch

configuring, 1-1

NetWorker

update, 1-7

NFS

duplicate request cache scalability limitations, 2-2
tuning the duplicate request cache, 2-4

P

Packetfilter

dbx restriction, 1-4
enhancements, 1-4

Pascal

privileged program failure, 1-3

R

Retirement

Aurema ARMTech Products, 1-8
DEC Ada, 1-9

S

Secure Web Server

displaying Apache documentation, 2-1
sources, 1-8
update, 1-7

security

buffer overflow exploitation, 1-1

Security

javaexecutedata, 1-1

Sources for Open Source

components

contents, 1-8

T

tunable Big Pages attributes, 1-1

U

Unicensus RCM

update, 1-7

Uninstalling a patch kit

problems encountered after
changing the system
configuration, 2-5

V

Visual Threads

update, 1-7