



Solaris 10 What's New



Sun Microsystems, Inc.
4150 Network Circle
Santa Clara, CA 95054
U.S.A.

Part No: 817-0547-30
October 2009

Copyright 2009 Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 U.S.A. All rights reserved.

Sun Microsystems, Inc. has intellectual property rights relating to technology embodied in the product that is described in this document. In particular, and without limitation, these intellectual property rights may include one or more U.S. patents or pending patent applications in the U.S. and in other countries.

U.S. Government Rights – Commercial software. Government users are subject to the Sun Microsystems, Inc. standard license agreement and applicable provisions of the FAR and its supplements.

This distribution may include materials developed by third parties.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the U.S. and other countries, exclusively licensed through X/Open Company, Ltd.

Sun, Sun Microsystems, the Sun logo, the Solaris logo, the Java Coffee Cup logo, docs.sun.com, Java, and Solaris are trademarks or registered trademarks of Sun Microsystems, Inc. or its subsidiaries in the U.S. and other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc. FireWire is a trademark of Apple Computer, Inc., used under license. Netscape and Netscape Navigator are trademarks or registered trademarks of Netscape Communications Corporation. Mozilla is a trademark or registered trademark of Netscape Communications Corporation in the United States and other countries. PostScript is a trademark or registered trademark of Adobe Systems, Incorporated, which may be registered in certain jurisdictions. OpenGL is a registered trademark of Silicon Graphics, Inc.

The OPEN LOOK and Sun™ Graphical User Interface was developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a non-exclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUIs and otherwise comply with Sun's written license agreements.

Products covered by and information contained in this publication are controlled by U.S. Export Control laws and may be subject to the export or import laws in other countries. Nuclear, missile, chemical or biological weapons or nuclear maritime end uses or end users, whether direct or indirect, are strictly prohibited. Export or reexport to countries subject to U.S. embargo or to entities identified on U.S. export exclusion lists, including, but not limited to, the denied persons and specially designated nationals lists is strictly prohibited.

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

Copyright 2009 Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 U.S.A. Tous droits réservés.

Sun Microsystems, Inc. détient les droits de propriété intellectuelle relatifs à la technologie incorporée dans le produit qui est décrit dans ce document. En particulier, et ce sans limitation, ces droits de propriété intellectuelle peuvent inclure un ou plusieurs brevets américains ou des applications de brevet en attente aux Etats-Unis et dans d'autres pays.

Cette distribution peut comprendre des composants développés par des tierces personnes.

Certains composants de ce produit peuvent être dérivées du logiciel Berkeley BSD, licenciés par l'Université de Californie. UNIX est une marque déposée aux Etats-Unis et dans d'autres pays; elle est licenciée exclusivement par X/Open Company, Ltd.

Sun, Sun Microsystems, le logo Sun, le logo Solaris, le logo Java Coffee Cup, docs.sun.com, Java et Solaris sont des marques de fabrique ou des marques déposées de Sun Microsystems, Inc., ou ses filiales, aux Etats-Unis et dans d'autres pays. Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux Etats-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc. FireWire est une marque de Applex Computer, Inc., utilisé sous le permis. Netscape est une marque de Netscape Communications Corporation. Netscape Navigator est une marque de Netscape Communications Corporation. Mozilla est une marque de Netscape Communications Corporation aux Etats-Unis et à d'autres pays. PostScript est une marque de fabrique d'Adobe Systems, Incorporated, laquelle pourrait être déposée dans certaines juridictions. OpenGL est une marque d'pos?e de Silicon Graphics, Inc.

L'interface d'utilisation graphique OPEN LOOK et Sun a été développée par Sun Microsystems, Inc. pour ses utilisateurs et licenciés. Sun reconnaît les efforts de pionniers de Xerox pour la recherche et le développement du concept des interfaces d'utilisation visuelle ou graphique pour l'industrie de l'informatique. Sun détient une licence non exclusive de Xerox sur l'interface d'utilisation graphique Xerox, cette licence couvrant également les licenciés de Sun qui mettent en place l'interface d'utilisation graphique OPEN LOOK et qui, en outre, se conforment aux licences écrites de Sun.

Les produits qui font l'objet de cette publication et les informations qu'il contient sont régis par la législation américaine en matière de contrôle des exportations et peuvent être soumis au droit d'autres pays dans le domaine des exportations et importations. Les utilisations finales, ou utilisateurs finaux, pour des armes nucléaires, des missiles, des armes chimiques ou biologiques ou pour le nucléaire maritime, directement ou indirectement, sont strictement interdites. Les exportations ou réexportations vers des pays sous embargo des Etats-Unis, ou vers des entités figurant sur les listes d'exclusion d'exportation américaines, y compris, mais de manière non exclusive, la liste de personnes qui font objet d'un ordre de ne pas participer, d'une façon directe ou indirecte, aux exportations des produits ou des services qui sont régis par la législation américaine en matière de contrôle des exportations et la liste de ressortissants spécifiquement désignés, sont rigoureusement interdites.

LA DOCUMENTATION EST FOURNIE "EN L'ETAT" ET TOUTES AUTRES CONDITIONS, DECLARATIONS ET GARANTIES EXPRESSES OU TACITES SONT FORMELLEMENT EXCLUES, DANS LA MESURE AUTORISEE PAR LA LOI APPLICABLE, Y COMPRIS NOTAMMENT TOUTE GARANTIE IMPLICITE RELATIVE A LA QUALITE MARCHANDE, A L'APTITUDE A UNE UTILISATION PARTICULIERE OU A L'ABSENCE DE CONTREFACON.

Contents

Preface	23
1 What's New in the Solaris 10 10/09 Release	27
System Administration Enhancements	27
Two-Terabyte Disk Support for Installing and Booting the Solaris OS	27
pci tool Utility	28
ZFS Features and Changes	28
nss_ldap shadowAccount Support	42
Sun Validation Test Suite 7.0 Patch Set 6	42
Installation Enhancements	42
Turbo-Charging SVR4 Packaging	43
System Resources Enhancements	43
Zones Parallel Patching	43
PVIO Drivers in Solaris 10 Guest Domains	43
Device Management Enhancements	43
iSCSI Initiator SMF Service	43
LSI 6180 Controller Support in Solaris MPxIO	44
System Performance Enhancements	44
Callout Subsystem Scalability	44
Driver Enhancements	44
Enhanced 10 Gigabit Ethernet Performance	44
InfiniBand Feature Improvements	45
Driver for LSI MPT 2.0-Compliant SAS2.0 Controllers	46
x86: Broadcom NetXtreme II Gigabit Ethernet Driver	46
x86: Interrupt Remapping Support for Intel Vt-d	46
x86: SATA Tape Devices Supported by the AHCI Driver	46
Sun StorageTek 6Gb/s SAS PCIe RAID HBA Driver	46
Intel 82599 10Gb PCI Express Ethernet Controller	47

Intel 82598 10Gb PCI Express Ethernet Controller	47
Freeware Enhancements	47
NTP Version 4.2.5	47
PostgreSQL	47
Samba	47
2 What's New in the Solaris 10 5/09 Release	49
System Resources Enhancements	49
Support Added for Using ZFS Clones When Cloning a Zone	49
zoneadm attach -b Option	49
System Administration Enhancements	50
SMF Services for IPsec	50
Security Enhancements	50
NAT-Traversal for IPsec Key Management Developers	50
Stronger Algorithms for IPsec	51
SunSSH With OpenSSL PKCS#11 Engine Support	51
Device Management Enhancements	51
x86: T-State Support for Intel Based Processors	52
System Performance Enhancements	52
Large Segment Offload Support for Intel PCI Express 10Gb NIC Driver	52
Solaris Power Aware Dispatcher and Deep C-State Support	52
Developer Tools Enhancements	53
SunVTS 7.0 Patch Set 5	53
x86: CPU Performance Counter Updates for Intel Processors	54
Driver Enhancements	54
hermon Driver	54
iSCSI Target	54
x86: NetXen 10‐GigE Device Driver	55
Intel ICH10 and Hartwell NIC Support in E1000g Driver	55
xge Driver Can Enable Multiple Receive Rings and MSI-X	55
Language Support Enhancements	55
New Locale Support for Kazakhstan and Ukraine	55
Additional Software Enhancements	55
SPARC: Fp-scrubber Daemon	55

3 What's New in the Solaris 10 10/08 Release	57
Installation Enhancements	57
Solaris Installation for ZFS Root Pools	57
System Administration Enhancements	58
Solaris Installation Tool Support of ZFS File Systems	58
SunVTS 7.0 Patch Set 3	59
lockstat Provider for DTrace	59
System Resource Enhancements	59
New Solaris Zones Features	59
x86: New GRUB findroot Command	60
x64: Support for 256 Processors	60
System Performance Enhancements	61
SPARC: Solaris SPARC Boot Architecture Redesigned	61
x86: Kernel Support for Intel SSSE3, SSE4.1, SSE4.2, and AMD SSE4A	62
Security Enhancements	62
Separation of Duty Enforcement Through the Solaris Management Console	62
SHA256/SHA512 crypt(3C) Plug-in	62
pam_list Module	62
Desktop Enhancements	63
SPARC: Adobe Reader 8.1.2	63
Flash Player 9.0.124.0	63
Networking Enhancements	63
Communication Protocol Parser Utilities	63
SIP End-to-end Traffic Measurements and Logging	64
Device Management Enhancements	64
Faulty Device Retirement Feature	64
MPxIO Support for Hitachi Adaptable Modular Storage Series Arrays	65
Driver Enhancements	66
x86: NVIDIA ck804/mcp55 SATA Controller Driver	66
x86: LSI MegaRAID SAS Controllers Driver	66
ixgbe Driver	66
SPARC: Support for aac Driver	67
Additional Software Enhancements	67
Perl Database Interface and Perl PostgreSQL Driver	67
PostgreSQL 8.3	67
Language Support Enhancements	67

IIIMF Hangul Language Engine	67
Freeware Enhancements	67
C-URL - The C-URL Wrappers Library	68
Libidn - Internationalized Domain Library	68
LibGD - The Graphics Draw Library	68
TIDY HTML Library	68
4 What's New in the Solaris 10 5/08 Release	69
System Administration Enhancements	69
Solaris Trusted Extensions Administrator Procedures	69
Flash Update Tool	70
PPD File Management Utility	70
Internet Printing Protocol Client-Side Support	71
Selectable Use of <code>localhost</code> for Solaris Print Server Database Hostname	71
Fault Management for T5140/T5240 Platforms	72
SunVTS 7.0	72
System Resource Enhancements	73
Solaris Zones and Branded Zones	73
CPU Caps	73
<code>projmod(1M)</code> Option	73
Device Management Enhancements	74
Tape Self-Identification	74
x86: Enhanced Speedstep CPU Power Management	74
x86: PowerNow! CPU Performance Management	74
iSNS Support in the Solaris iSCSI Target	74
Security Enhancements	75
Solaris Trusted Extensions Supports Mounting Labeled File Systems With the NFSv3 Protocol	75
SPARC: Hardware -Accelerated Elliptical Curve Cryptography (ECC) Support	75
Networking Enhancements	75
Sockets Direct Protocol	75
<code>inetd</code> Backlog Queue Size	76
X11 Windowing Enhancements	76
Xvnc Server and Vncviewer Client	76
Desktop Tools Enhancements	76

StarOffice 8	77
Flash Player 9	77
Pidgin 2.0	77
PAPI Print Commands	77
System Performance Enhancements	78
64-bit SPARC: Memory Placement Optimization Support for sun4v Platforms	78
SPARC: Shared Contexts Support	78
x86: CPUID-Based Cache Hierarchy Awareness	79
Language Support Enhancements	79
Locale Creator	79
libchewing 0.3.0	79
File Encoding Examiner	80
Kernel Functions Enhancements	80
x86: MONITOR and MWAIT CPU Idle Loop	80
Driver Enhancements	80
x86: Support Sun Fire X4540 Disk Status Indicators	81
MPxIO Extension for Serial Attached SCSI Devices on mpt(7D)	81
x86: SATA ATAPI Support in AHCI Driver	81
x86: AMD-8111	81
SATA NCQ Support in AHCI Driver	81
x86: bnx II Ethernet Driver	81
USB-to-Serial Driver for Keyspan Adapters	82
Freeware Enhancements	82
32-bit: pgAdmin III	82
p7zip	82
5 What's New in the Solaris 10 8/07 Release	83
System Administration Enhancements	83
Name Service Switch Enhancements	83
iostat Improvements	83
Solaris System Registration	83
Sun Service Tag	84
MPxIO Path Steering	84
raidctl	85
Brand-Specific Handlers for zoneadm Commands	85

x86: Fault Management For Next Generation AMD Opteron Processors	85
x86: Predictive Self-Healing for PCI Express on x64 Systems	86
x86: stmsboot Porting	86
x86: Concurrent FPDMA READ/WRITE QUEUED Under SATA Module	86
x86: Tagged Queuing	86
Installation Enhancements	86
NFSv4 Domain Name Configurable During Installation	86
Solaris Live Upgrade	87
Upgrading the Solaris OS When Non-Global Zones Are Installed	87
Keyboard Configuration Automated	89
Deferred-Activation Patching	89
Networking Enhancements	90
IPsec Tunnel Reform	90
Packet Filter Hooks	91
SMF Enhancements to Routing Management	91
Quagga Software Routing Suite	91
DHCPv6 Client	91
Single Hosts File	92
Large Send Offload	92
x86: nge Driver Updated to Support Jumbo Framework	92
NFSv4 Domain Name Configurable During Installation	92
Security Enhancements	92
Solaris Key Management Framework	92
libmd - Message Digest library	93
Solaris Cryptographic Framework	93
Solaris Data Encryption Supplement	93
File System Enhancements	93
Support for iSCSI Target Devices	94
Extended FILE Space for 32-bit Solaris Processes	94
System Resources Enhancements	94
lx Branded Zones: Solaris Containers for Linux Applications	94
Improved zonecfg Procedures for Creating Containers	95
IP Instances: LAN and VLAN Separation for Non-Global Zones	96
Solaris Zones Boot Enhancements	96
System V Resource Controls for Zones	97
Zone Unique Identifier	98

Ability to Mark Zones as “Incomplete”	98
Using DTrace in a Non-Global Zone	98
Desktop Tools Enhancements	98
Thunderbird 2.0	99
Firefox 2.0 Web Browser	99
Gaim OTR plugin	99
x86: XVideo Support RealPlayer	99
X11 Windowing Enhancements	99
dtlogin Language Selection Overhaul	99
X Server DTrace Provider	99
Xorg X11R7.2 Server and Drivers	100
Language Support Enhancements	100
Migration of Existing EMEA, Central and South American Locales to Common Locale Data Repository	100
Japanese Font Update	100
More Japanese iconv Modules for Unicode	100
Input Method Switcher Enhancement and EMEA Keyboard Layout Emulation Support	101
x86: Zero-CountryCode Keyboard Layout Support	101
Developer Tool Enhancements	101
SunVTS 6.4	101
Driver Enhancements	102
Reliable Datagram Sockets	102
Enhanced USB EHCI Host Controller Driver	102
USCSI LUN Reset Support	102
SATA HBA Framework and Marvell Driver	102
Compact Flash Support	103
ACM driver of USB Communication Device Class	103
CardBus Support	103
IBM LTO-4 Tape Drive Support	103
HP LTO-4 Tape Drive Support	103
NVIDIA Accelerated Graphics Drivers	103
SPARC: ntwdt Driver for UltraSPARC-T1 (Niagara) Systems	103
x86: ACPI Thermal Zone Monitor	103
x86: Adaptec aac Hardware Support	104
x86: Solaris Audio Driver for ATI IXP400	104

x86: High-Definition Audio Driver	104
x86: SATA AHCI HBA Driver	104
System Performance Enhancements	104
SPARC: UltraSPARC T2 PCI Express Interface Unit Performance Counter Data	104
Hashed Cache Index Mode Support	105
Multi-level CMT Scheduling Optimizations	105
Process Count Scalability	105
MPSS Extended to Shared Memory	106
Device Management Enhancements	106
Enhanced st SCSI Reservations	106
CPU Power Management	106
Console Subsystem Enhancements	107
Coherent Console	107
6 What's New in the Solaris 10 11/06 Release	109
System Administration Enhancements	109
Storage Networking Industry Association Multipath Management API Support	109
Sun Java Web Console Changes	109
File-System Monitoring Tool	110
System Resource Enhancements	110
Resource Management Features	110
Solaris Zones Features	111
Logical Domains Features	113
Security Enhancements	113
Solaris Trusted Extensions	113
Solaris Trusted Extensions for Printing	114
Solaris Trusted Extensions File-System Labeling	114
Device Management Enhancements	114
Support for PCI Express (PCIe)	114
x86: Sun Fire X4500 SATA Disk FMA	115
SPARC: Transitioning SPARC-Based Systems From Ipge to E1000g Network Drivers	115
Solaris Fibre Channel Host-Based Logical Unit Number Masking	116
SPARC: Extended Message Signaled Interrupt Support for Fire-Based Platforms	116
Improved Device in Use Error Checking	116
Desktop Enhancements	117

Default Desktop Session in dt login	117
Adobe Flash Player Plugin for Solaris	117
GNOME-VFS and Nautilus ACL Support	117
Solaris Trusted Extensions Desktops	117
Installation Enhancements	118
Solaris Flash Archives	118
Secure By Default Network Profile	118
Installing Solaris Trusted Extensions	119
System Performance Enhancements	120
SPARC: Watchdog Timer for Sun4V	120
Networking Enhancements	120
Sun Java System Message Queue 3.7 Update 1	120
New and Updated Drivers	120
ST Driver Support for Quantum LTO-2 and LTO-3 Tape Drives	120
CDB Length Capability	120
Language Support	120
IIIMF and Language Engines	121
7 What's New in the Solaris 10 6/06 Release	123
System Administration Enhancements	123
The Solaris ZFS File System	123
Predictive Self-Healing for x64 Systems	126
Predictive Self-Healing Support for SNMP Notification	127
SunVTS 6.2	127
Common Agent Container	128
iSCSI Logout Support	128
iSCSI MS/T Support	128
The logadm Utility	129
The volfsUtility	129
Basic Registration 1.1	129
Sun Update Connection	130
Networking Enhancements	130
IPv6 for IPFilter	130
UDP and TCP Performance Enhancement	130
IP_NEXTHOP Socket Option	131

TCP_INIT_CWND Socket Option	131
Security Enhancements	131
Enhancements for pkttool Object Migration and Interoperability	131
SSL Proxy Module	132
AES Counter Mode	132
PKCS #11 v2.20 Support in the Solaris Cryptographic Framework	132
Kerberos Cred Auto-Renew	132
Device Management Enhancements	132
iSNS Client Support for iSCSI	133
cdrecord, readCD, and cdda2wav Available	133
x86: PCI Express Support on x86 Systems	133
Solaris support for LSISAS1064 RAID Operations	134
Desktop Enhancements	135
32-bit: Palm Synchronization with USB Port	135
32-bit: The gnome-pilot Utility	135
x86: Xorg X Server Version 6.9	135
Developer Tools Enhancements	135
mediaLib 2.4	135
New and Updated Drivers	136
x86: SATA HBA Framework Support	136
USB-to-Serial Driver for Prolific Adapters	136
Driver for IEEE 1394-Based (IIDC) Digital Cameras	136
Driver Support for the New STK Tape Drive 10000 “Titanium”	136
USB-to-Serial Driver for Keyspan Adapters	137
Deimos Cryptographic Accelerator	137
x86: Driver Support for AMD64 Platform	137
rge Driver	137
Chelsio NIC Driver Support	138
HBA Drivers	138
Language Support Enhancements	138
New UTF-8 Locales	138
Additional Software	138
Watchdog Timeout	138
32-bit: RealPlayer for the Solaris OS	139
pilot-link Software	139
PostgreSQL for the Solaris OS	139

8 What's New in the Solaris 10 1/06 Release	141
Installation Enhancements	141
Upgrade Support Changes for Solaris Releases	141
Sun Update Connection, System Edition 1.0	141
x86: GRUB Based Booting	142
Upgrading the Solaris OS When Non-Global Zones Are Installed	143
Networking Enhancements	144
Source-Filtered Multicasting	144
Enhancements for bge and xge Network Interfaces	144
Enhancements to Java Desktop System Release 3	145
New Features in Java DS	145
Plug-ins for Mozilla 1.7	146
Localization Support	147
Security Enhancements	147
Set SMTP to Use Transport Layer Security	147
Metaslot in the Cryptographic Framework	148
IKE Enhancements	148
New Command for embedded_su	148
System Performance Enhancements	149
Large Pages Enhancement	149
Kernel Page Relocation	149
Memory Placement Optimization Hierarchical Lgroup Support	149
System Administration Enhancements	149
Improvements to Volume Management (vol)	149
SunVTS 6.1	150
Support for SCSI Disks Larger Than 2 Terabytes	150
Fibre-Channel HBA Port Utility	151
Additional Banner Page Printing Option in Solaris Print Manager	151
x86: New prtconf Option to Display Product Names	151
Developer Enhancements	152
mediaLib 2.3	152
Advanced DDI Interrupts	152
Desktop Enhancements	153
SPARC: Adobe Acrobat Reader 7.0.1	153
SPARC: Sun OpenGL 1.5 for Solaris	154
Xorg Configuration GUI	154

Xorg Release 6.8.2	154
X Client Support for XFree86 Extensions	154
Support for Virtual USB Keyboard and Mouse Devices	155
Device Management Enhancements	155
Support for iSCSI Devices	155
New and Updated Drivers	156
SPARC: Sun XVR-2500 Graphics Accelerator	156
SCSI HBA Driver for LSI MegaRAID 320-2x SCSI RAID Controller	156
USB CCID IFD Handler	156
New nge Driver	157
Documentation Enhancements	157
Solaris 10 1/06 Documentation	157
Sun Fire Documentation Changes	157
9 What's New in the Solaris 10 3/05 Release	159
Key Features in Solaris 10 Software	159
System Administration Enhancements	160
Predictive Self-Healing	160
DTrace Dynamic Tracing Facility	162
64-bit Support on x86 Systems	163
Kernel Selection for x86 Systems	163
Sun Java Web Console	163
Multi-Owner Disk Set Support With Solaris Volume Manager for Sun Cluster	164
Expanded Disk Set Support in the Solaris Volume Manager	164
Importing Remotely Replicated Disk Sets With Solaris Volume Manager	165
Device ID Changes	165
Multiterabyte Volume Support in Solaris Volume Manager	165
Reconfiguration Coordination Manager Support in Solaris Volume Manager	166
Top-Down Volume Creation in Solaris Volume Manager	166
Package and Patch Tool Enhancements	166
Enhancements for <code>pbind</code> and <code>psrset</code> Commands	167
Solaris Dynamic Host Configuration Protocol (DHCP)	167
DHCP Event Scripting	167
DHCP for Logical Interfaces	168
x86: SunVTS 6.0	168

Kernel Modular Debugger	169
Solaris Process Accounting and Statistics Improvements	169
Enhancements to <code>ls</code> Command	170
New <code>ps rinfo</code> Option to Identify Chip Multithreading Features	170
Enhancements to <code>pfiles</code> Tool	170
Solaris IP Filter	170
Core File Content Enhancements	171
System Management Agent	172
SPARC: 64-bit Package Changes for System Administrators	172
NIS-to-LDAP Transition Service	172
Signed Packages and Patches	173
System V IPC Configuration	173
Optional Interval for <code>netstat</code> Command	175
GMT-offset Quoted Time Zones	175
Combine Data Address and Test Address for Single IP Network Multipathing Group	175
LDAP Command Changes	175
Desktop Enhancements	176
Java Desktop System, Release 3	176
Mozilla 1.7	178
System Resources Enhancements	178
Solaris Zones Software Partitioning Technology	178
System V IPC and Other Resource Controls	180
New Solaris Project and Resource Management Command Functionality	181
Dynamic Resource Pools	181
Extended Accounting Subsystem Enhancements	182
Physical Memory Control Using the Resource Capping Daemon	183
Enhancements to the Java 2 Platform, Standard Edition 5	183
New Additions to the Java Programming Language	184
Monitoring and Manageability Enhancements for Java Platform	184
Performance and Scalability Improvements for Java Platform	184
XML 1.1 and Namespace, XSLTC, SAX 2.0.2, and DOM Level 3, Java based APIs	185
New Default Swing Look and Feel	185
Diagnostic Tools	185
Unicode Support	185
Installation Enhancements	186
Solaris Installation Changes Including Installation Unification	186

Custom JumpStart Installation Package and Patch Enhancements	187
x86: Specifying Boot Properties With <code>add_install_client</code> Command	188
Configuring Multiple Network Interfaces During Installation	188
Replacement of Solaris Installation Command-Line Interface	189
SPARC: 64-bit Package Changes	189
WAN Boot Installation Method	189
Creating RAID-1 Volumes (Mirrors) With Custom JumpStart	190
Solaris Live Upgrade 2.1	190
Custom JumpStart Installation Method Creates New Boot Environment	191
Solaris Flash Archives	191
Solaris Flash Differential Archives and Configuration Scripts	191
Customizing Content in a Solaris Flash Archive	192
Command-Line Interface Enhancements to the Solaris Product Registry	193
Reduced Networking Software Group	193
Modifying Disk Partition Tables by Using a Virtual Table of Contents (VTOC)	194
x86: Change in Default Boot-Disk Partition Layout	194
Support for Lightweight Directory Access Protocol (LDAP) Version 2 Profiles	194
Security Enhancements	195
Signing ELF Objects	195
Process Rights Management	195
Changes to PAM for the Solaris 10 OS	196
<code>pam_ldap</code> Changes	197
Solaris Secure Shell Enhancements	198
OpenSSL and OpenSSL PKCS#11 Engine	199
<code>sshd</code> Daemon and <code>/etc/default/login</code>	199
New Password Options for Nonlogin and Locked Accounts	199
<code>-setcond</code> Option to <code>auditconfig</code> Command Is Removed	200
<code>perzone</code> Audit Policy	200
Kerberos Enhancements	200
TCP Wrappers for <code>rpcbind</code>	202
<code>zonename</code> Audit Token and Audit Policy Option	202
User Commands for Solaris Cryptographic Framework	202
IKE Configuration Parameters	203
Simple Authentication and Security Layer	203
Audit Time Now Reported in ISO 8601 Format	203
Basic Audit and Reporting Tool	204

IPsec and the Solaris Cryptographic Framework	204
Solaris Cryptographic Framework for System Administrators	205
Remote Audit Log	205
FTP Server Enhancements	206
FTP Client	207
Internet Key Exchange (IKE) Key Storage on Sun Crypto Accelerator 4000 Board	207
IKE Hardware Acceleration	208
ipseckey Enhancement	208
Credential Propagation Over Loopback Connections	208
Auditing Header Token Contains Host Information	209
Auditing Enhancements	209
New Audit Token path_attr	210
Password History Checking	210
Enhanced crypt () Function	211
System Performance Enhancements	211
New Architecture for Network Stacks	211
CPU Performance Counters	211
Improved System Performance for Large Numbers of Interfaces	212
UFS Logging Performance Enhancements	212
Memory Placement Optimization (MPO)	212
Dynamic Intimate Shared Memory (DISM) Large-Page Support	213
Device Management	213
New or Updated Drivers	213
1394 (FireWire) and Mass-Storage Support on x86 Systems	215
Internet Printing Protocol Listener	216
Fibre Channel Connectivity for Storage Devices	216
Expanded Printer Support	216
Common Solaris Target Disk Driver	217
Wheel Mouse Support	217
USB 2.0 Features	218
USB 2.0 Devices	218
Solaris Support for USB Devices	219
USB Mass Storage Devices	219
USB Driver Enhancements	220
EHCI and OHCI Drivers	221
Logical Unit Reset	222

Networking Enhancements	222
Virtual IP Source Address Selection	222
Stream Control Transmission Protocol	223
Zebra Multiprotocol Routing Suite	223
IPsec and NAT Traversal	223
Enhancement to the nfsmapid Daemon	224
sendmail Version 8.13	225
sendmail Version 8.12 Uses TCP Wrappers	225
Sun Java System Message Queue	226
Sun Java System Application Server	227
Using CacheFS With NFS Version 4	227
Enhancement for vacation Utility	228
MILTER, sendmail's New Mail Filter API	228
IPv6 Advanced Sockets API	229
Contents of /usr/lib/mail Have Moved to /etc/mail/cf	229
Added IPv6 Functionality During Solaris Installation	229
IPv6 Temporary Addresses	230
routeadm Command	230
TCP Multidata Transmit	230
ifconfig router Option	231
Internet Protocol Version 6 (IPv6) Default Address Selection	231
Disable NFS and automount Services	232
Internet Protocol Version 6 (IPv6) 6to4 Router	233
Packet Tunneling Over IPv6	233
Hosting Multiple Web Sites on a Single Solaris Machine	233
IP Quality of Service	234
User Selector for Internet Protocol Quality of Service (IPQoS)	234
Routing Information Protocol version 2 (RIPv2)	234
Developer Tool Enhancements	235
Dynamic Tracing Facility	235
GCC Version 3.4.3	235
Perl Version 5.8.4	235
Per-thread Mode Enhancement	236
USB End-User Device Support Enhancements	236
ls Enhancements	236
New Functions for Converting Strings	236

Java Support for pstack Command	237
New Mechanisms for the Solaris Cryptographic Framework	237
Retail and Nonretail Options for Providers in Solaris Cryptographic Framework	237
Linkers and Libraries Updates	238
Layered Driver Interfaces	239
Changes to makecontext () Function	239
Single UNIX Specification, Version 3	240
Advanced API	240
Simple Authentication and Security Layer for Developers	240
Event Ports	241
Core File Content	241
Atomic Operations	241
Solaris WBEM File Changes	241
Privileges for Software Developers	242
Solaris Cryptographic Framework for Developers	242
SPARC: 64-bit Package Changes for Software Developers	243
SPNEGO Pseudo-Mechanism for GSS-API Applications	244
Locality Groups	244
Thread Stacks in pmap	244
New DOOR_REFUSE_DESC Flag	244
Stack Check APIs	245
Enhanced crypt () Function for Software Developers	245
New Flags for madvise () Function	246
Memory Allocation With libumem	246
Smart Card Terminal Interfaces	246
Smartcard Middleware APIs	246
File System Enhancements	247
NFS Version 4	247
UFS Logging Enabled by Default	247
Enhancements to the NFS Client	248
Multiterabyte UFS File Systems	249
Device File System (devfs)	250
Multiterabyte Disk Support With EFI Disk Labels	250
New Configuration File for the Autofs Environment	251
X11 Windowing Enhancements	252
Xorg X Server	252

Xfixes Xserver Extension	252
Xscreensaver Program	252
X Event Interception Extension (XEvIE)	253
FreeType 2.1.x	253
Xserver Virtual Screen Enhancement	253
Xrender Extension	253
Language Support Enhancements	254
Unicode Version 4.0 Support	254
Code Conversions for Internationalized Domain Name Support	254
New iconv Code Conversions	255
New Solaris Unicode Locales	255
Locale Administrator	255
Standard Type Services Framework	256
Auto Encoding Finder	256
SunIM-based Korean Language Engine With Auxiliary Window Support	256
Common Transliteration-based Input Method for All Indian Languages	257
Additional Keyboard Support	257
Wubi Input Method	257
Input Method Support for Indic	258
Seven Additional Indic Scripts for Support in Unicode Locales	258
HKSCS-2001 Support in Hong Kong Locales	258
Additional Software	259
Sun Remote Services Net Connect 3.1.1	259
Freeware Enhancements	259
GCC v.3.4.3	259
Webmin Tool	259
Intelligent Platform Management Interface	259
Apache Version 2	260
BIND 9	260
Samba Enhancements	260
Flex 2.5.4a	261
SIP Proxy Server	261
libusb 0.1.8	262
Ghostscript 7.05	262
New Freeware Packages: libxml2 2.4.16 and libxslt 1.0.19	262
ANT 1.4.1 Freeware Package	262

Documentation Changes	263
Solaris 10 Documentation DVD	263
New Documents	263
Document Reorganizations	265
Documentation That Has Been Moved	267
Documentation That Is Not Published in This Release	267
Features by Software Express Release Date	267
New Features in Solaris 10 1/06 Release	267
New Features in Solaris 10 3/05 Release	269

Preface

Solaris 10 What's New summarizes all features in the Solaris™ 10 OS that are new or have been enhanced since the Solaris 9 Operating System (Solaris OS) was originally distributed in May 2002. The current release is the Solaris 10 10/09 release.

Note – For more information about features in the current release, see *Solaris 10 10/09 What's New*.

In the Solaris 10 OS, Sun Microsystems has developed a new architecture for building and deploying systems and services that are capable of Predictive Self-Healing. See “[Predictive Self-Healing](#)” on page 160. Also, changes in the installation of the Solaris OS provide a simplified and unified installation process. See “[Solaris Installation Changes Including Installation Unification](#)” on page 186.

The Solaris Zones partitioning technology is a key feature that has been added to the Solaris 10 OS. Zones are used to virtualize operating system services and provide an isolated and secure environment for running applications. See “[Solaris Zones Software Partitioning Technology](#)” on page 178. Additional key features in the Solaris 10 OS include “[Process Rights Management](#)” on page 195, “[DTrace Dynamic Tracing Facility](#)” on page 162, and “[New Architecture for Network Stacks](#)” on page 211. “[Enhancements to Java Desktop System Release 3](#)” on page 145 are now also available in the Solaris 10 OS.

Note – This Solaris release supports systems that use the SPARC® and x86 families of processor architectures. The supported systems appear in the [Solaris OS: Hardware Compatibility Lists](#) (<http://www.sun.com/bigadmin/hcl>). This document cites any implementation differences between the platform types.

In this document these x86 related terms mean the following:

- “x86” refers to the larger family of 64-bit and 32-bit x86 compatible products.
- “x64” relates specifically to 64-bit x86 compatible CPUs.
- “32-bit x86” points out specific 32-bit information about x86 based systems.

For supported systems, see the *Solaris OS: Hardware Compatibility Lists*.

Who Should Use This Book

This book provides introductory descriptions of the new Solaris 10 features for users, developers, and system administrators who install and use the Solaris 10 Operating System.

Optional Feature Licensing

Certain optional features and products that are described in this document might require individual licensing for use. Refer to the Software License Agreement.

Related Books

For further information about the features that are summarized in this book, refer to the following Solaris 10 documentation at <http://docs.sun.com/app/docs/prod/solaris.10>.

Related Third-Party Web Site References

Third-party URLs are referenced in this document and provide additional, related information.

Note – Sun is not responsible for the availability of third-party Web sites mentioned in this document. Sun does not endorse and is not responsible or liable for any content, advertising, products, or other materials that are available on or through such sites or resources. Sun will not be responsible or liable for any actual or alleged damage or loss caused by or in connection with use of or reliance on any such content, goods, or services that are available on or through such sites or resources.

Documentation, Support, and Training

The Sun web site provides information about the following additional resources:

- [Documentation \(http://www.sun.com/documentation/\)](http://www.sun.com/documentation/)
- [Support \(http://www.sun.com/support/\)](http://www.sun.com/support/)
- [Training \(http://www.sun.com/training/\)](http://www.sun.com/training/)

Sun Welcomes Your Comments

Sun is interested in improving its documentation and welcomes your comments and suggestions. To share your comments, go to <http://docs.sun.com> and click Feedback.

Typographic Conventions

The following table describes the typographic conventions that are used in this book.

TABLE P-1 Typographic Conventions

Typeface	Meaning	Example
AaBbCc123	The names of commands, files, and directories, and onscreen computer output	Edit your <code>.login</code> file. Use <code>ls -a</code> to list all files. <code>machine_name% you have mail.</code>
AaBbCc123	What you type, contrasted with onscreen computer output	<code>machine_name% su</code> Password:
<i>aabbcc123</i>	Placeholder: replace with a real name or value	The command to remove a file is <i>rm filename</i> .
<i>AaBbCc123</i>	Book titles, new terms, and terms to be emphasized	Read Chapter 6 in the <i>User's Guide</i> . A <i>cache</i> is a copy that is stored locally. Do <i>not</i> save the file. Note: Some emphasized items appear bold online.

Shell Prompts in Command Examples

The following table shows the default UNIX® system prompt and superuser prompt for the C shell, Bourne shell, and Korn shell.

TABLE P-2 Shell Prompts

Shell	Prompt
C shell	<code>machine_name%</code>
C shell for superuser	<code>machine_name#</code>
Bourne shell and Korn shell	<code>\$</code>
Bourne shell and Korn shell for superuser	<code>#</code>

What's New in the Solaris 10 10/09 Release

This document summarizes all features in the Solaris 10 Operating System that are new or have been enhanced since the Solaris 9 OS was originally distributed in May 2002. This chapter summarizes new features in the current release, the Solaris 10 10/09 release. [Chapter 2, “What's New in the Solaris 10 5/09 Release,”](#) summarizes new features in the previous Solaris 10 5/09 release. [Chapter 3, “What's New in the Solaris 10 10/08 Release,”](#) summarizes new features in the Solaris 10 10/08 release. [Chapter 4, “What's New in the Solaris 10 5/08 Release,”](#) summarizes new features in the Solaris 10 5/08 release. [Chapter 5, “What's New in the Solaris 10 8/07 Release,”](#) summarizes new features in the Solaris 10 8/07 release. [Chapter 6, “What's New in the Solaris 10 11/06 Release,”](#) summarizes new features in the Solaris 10 11/06 release. [Chapter 7, “What's New in the Solaris 10 6/06 Release,”](#) summarizes new features in the Solaris 10 6/06 release. [Chapter 8, “What's New in the Solaris 10 1/06 Release,”](#) summarizes new features in the Solaris 10 1/06 release. [Chapter 9, “What's New in the Solaris 10 3/05 Release,”](#) summarizes new features in the Solaris 10 3/05 release. [Chapter 9, “What's New in the Solaris 10 3/05 Release,”](#) also summarizes all features, which are sorted by the Software Express release that introduced these features.

System Administration Enhancements

The following system administration features and enhancements have been added to the Solaris 10 10/09 release.

Two-Terabyte Disk Support for Installing and Booting the Solaris OS

Starting with the Solaris 10 10/09 release, you can install and boot the Solaris OS from a disk that is up to 2 Tbytes in size. In previous Solaris releases, you could not install and boot the Solaris OS from a disk that was greater than 1 Tbyte in size.

In this Solaris release, you can use the VTOC label on a disk of any size. However, the addressable space by the VTOC is limited to 2 Tbytes. This feature enables disks that are larger than 2 Tbytes to also be used as boot drive. However, the usable space from label is limited to 2 Tbytes.

This feature is available only on systems that run a 64-bit kernel. A minimum of 1 Gbyte of memory is required for x86-based systems.

For more information about the Solaris disk drivers and disk utilities that have been updated to support boot on disks greater than 1 Tbyte, see *System Administration Guide: Devices and File Systems*.

pci tool Utility

The `pci tool` utility enables system administrators to bind interrupts to specific hardware strands for enhanced performance. This utility exists in the public `SUNWio-tools` package. For more information about using `pci tool`, see the `pci tool` man page.

ZFS Features and Changes

The following section summarizes new features in the ZFS file system.

- **ZFS and Flash installation support** – In the Solaris 10 10/09 release, you can set up a JumpStart profile to identify a flash archive of a ZFS root pool. For more information, see the *Solaris ZFS Administration Guide*.
- **Setting ZFS user and group quotas** – In previous Solaris releases, you could apply quotas and reservations to ZFS file systems to manage and reserve space. In this Solaris release, you can set a quota on the amount of space consumed by files that are owned by a particular user or group. You might consider setting user and group quotas in an environment with a large number of users or groups. You can set user or group quotas by using the `zfs userspace` and `zfs groupspace` properties as follows:

```
# zfs set userquota@user1=5G tank/data
# zfs set groupquota@staff=10G tank/staff/admins
```

You can display a user's or group's current quota setting as follows:

```
# zfs get userquota@user1 tank/data
NAME          PROPERTY          VALUE          SOURCE
tank/data    userquota@user1  5G             local
# zfs get groupquota@staff tank/staff/admins
NAME          PROPERTY          VALUE          SOURCE
tank/staff/admins groupquota@staff  10G            local
```

- **Using ZFS ACL pass through inheritance for execute permission** – In previous Solaris releases, you could apply ACL inheritance so that all files are created with `0664` or `0666` permissions. If you want to optionally include the execute bit from the file creation mode into the inherited ACL, you can use the pass through inheritance for execute permission in this release.

If `aclinherit=passthrough-x` is enabled on a ZFS dataset, you can include execute permission for an output file that is generated from `cc` or `gcc` tools. If the inherited ACL does not include execute permission, then the executable output from the compiler won't be executable until you use the `chmod` command to change the file's permissions.

- **Using cache devices in your ZFS storage pool** – In the Solaris 10 10/09 release, you can create pool and specify *cache devices*, which are used to cache storage pool data. Cache devices provide an additional layer of caching between main memory and disk. Using cache devices provide the greatest performance improvement for random read-workloads of mostly static content.

One or more cache devices can specified when the pool is created. For example:

```
# zpool create pool mirror c0t2d0 c0t4d0 cache c0t0d0
# zpool status pool
  pool: pool
  state: ONLINE
  scrub: none requested
  config:
```

NAME	STATE	READ	WRITE	CKSUM
pool	ONLINE	0	0	0
mirror	ONLINE	0	0	0
c0t2d0	ONLINE	0	0	0
c0t4d0	ONLINE	0	0	0
cache				
c0t0d0	ONLINE	0	0	0

```
errors: No known data errors
```

For information about determining whether using cache devices is appropriate for your environment, see the [Solaris ZFS Administration Guide](#).

- **ZFS property enhancements** – The Solaris 10 10/09 release provides the following ZFS property enhancements:
 - You can set ZFS file system properties at pool creation time. In the following example, compression is enabled on the ZFS file system that is created when the pool is created.

```
# zpool create -O compression=on pool mirror c0t1d0 c0t2d0
```

- You can set two cache properties on a ZFS file system that allow you to control what is cached in the primary cache (ARC) or the secondary cache (L2ARC). The cache properties are set as follows:
 - `primarycache` – Controls what is cached in the ARC.
 - `secondarycache` – Controls what is cached in the L2ARC.

You can set these properties on an existing file system or when the file system is created. For example:

```
# zfs set primarycache=metadata tank/datab
# zfs create -o primarycache=metadata tank/newdatab
```

Some database environments might benefit from not caching user data. You will have to determine if setting cache properties is appropriate for your environment.

For more information, see the *Solaris ZFS Administration Guide*.

- You can use the space usage properties to identify space usage for clones, file systems, and volumes, but not snapshots. The properties are as follows:
 - `usedbychildren` – Identifies the amount of space that is used by children of this dataset, which would be freed if all the dataset's children were destroyed. The property abbreviation is `usedchild`.
 - `usedbydataset` – Identifies the amount of space that is used by this dataset itself, which would be freed if the dataset was destroyed, after first destroying any snapshots and removing any `refreservation`. The property abbreviation is `usedds`.
 - `usedbyrefreservation` – Identifies the amount of space that is used by a `refreservation` set on this dataset, which would be freed if the `refreservation` was removed. The property abbreviation is `usedrefreserv`.
 - `usedbysnapshots` – Identifies the amount of space that is consumed by snapshots of this dataset. In particular, it is the amount of space that would be freed if all of this dataset's snapshots were destroyed. Note that this is not simply the sum of the snapshots' `used` properties, because space can be shared by multiple snapshots. The property abbreviation is `usedsnap`.

These new properties break down the value of the `used` property into the various elements that consume space. In particular, the value of the `used` property breaks down as follows:

$$\text{used property} = \text{usedbychildren} + \text{usedbydataset} + \text{usedbyrefreservation} + \text{usedbysnapshots}$$

You can view these properties by using the `zfs list -o space` command. For example:

```
# zfs list -o space
NAME                AVAIL  USED  USEDSNAP  USEDDS  USEDREFRESERV  USEDCHILD
pool                33.2G  72K   0         21K     0              51K
rpool               27.0G  6.27G  20.5K    97K     0              6.27G
rpool/ROOT          27.0G  4.73G  0         21K     0              4.73G
rpool/ROOT/zfsBE    27.0G  4.73G  97.5M    4.63G   0              0
rpool/dump           27.0G  1.00G  16K      1.00G   0              0
rpool/export        27.0G  60K   16K      23K     0              21K
rpool/export/home   27.0G  21K   0         21K     0              0
rpool/swap           27.5G  553M  0         41.5M   512M          0
```

- In this release, snapshots are omitted from `zfs list` output. The `listsnapshots` pool property controls whether snapshot information is displayed by the `zfs list` command. If you use the `zfs list -t snapshots` command, snapshot information is displayed. The default value is `off`, which means snapshot information is not displayed by default.
- **ZFS log device recovery** – In the Solaris 10 10/09 release, ZFS identifies intent log failures in the `zpool status` command. FMA reports these errors as well. Both ZFS and FMA describe how to recover from an intent log failure.

For example, if the system shuts down abruptly before synchronous write operations are committed to a pool with a separate log device, you will see intent-log related error messages in the `zpool status` output. For information about resolving log device failures, see the [Solaris ZFS Administration Guide](#).

- **Using ZFS ACL Sets** – The Solaris 10 10/09 release provides the ability to apply NFSv4-style ACLs in sets, rather than apply different ACL permissions individually. The following ACL sets are provided:
 - `full_set` = all permissions
 - `modify_set` = all permissions except `write_acl` and `write_owner`
 - `read_set` = `read_data`, `read_attributes`, `read_xattr`, and `read_acl`
 - `write_set` = `write_data`, `append_data`, `write_attributes`, and `write_xattr`

These ACL sets are pre-defined and cannot be modified.

- **Zone migration in a ZFS environment** – In the Solaris 10 5/09 release, support for migrating zones in a ZFS environment with Live Upgrade is extended. For more information, see the [Solaris ZFS Administration Guide](#).
- **ZFS installation and boot support** – In the Solaris 10 10/08 release, you can install and boot a ZFS root file system. The initial installation option or the JumpStart feature is available to install a ZFS root file system. Or, you can use the Solaris Live Upgrade feature to migrate a UFS root file system to a ZFS root file system. ZFS support for swap and dump devices is also provided.
- **Rolling back a ZFS dataset without unmounting** – In the Solaris 10 10/08 release, you can rollback a dataset without unmounting it first. This feature means that the `zfs rollback -f` option is no longer needed to force an unmount operation. The `-f` option is no longer supported, and is ignored if specified.
- **Enhancements to the `zfs send` command** – The Solaris 10 10/08 release includes the following enhancements to the `zfs send` command:
 - You can send all incremental streams from one snapshot to a cumulative snapshot. For example:

```
# zfs list
NAME                                USED  AVAIL  REFER  MOUNTPOINT
pool                                428K  16.5G   20K    /pool
pool/fs                             71K   16.5G   21K    /pool/fs
```

```

pool/fs@snapA          16K    - 18.5K  -
pool/fs@snapB          17K    - 20K    -
pool/fs@snapC          17K    - 20.5K  -
pool/fs@snapD          0      - 21K    -
# zfs send -I pool/fs@snapA pool/fs@snapD > /snaps/fs@combo

```

This syntax shows how to send all incremental snapshots between fs@snapA to fs@snapD to fs@combo.

- You can send an incremental stream from the origin snapshot to create a clone. The original snapshot must already exist on the receiving side to accept the incremental stream. For example:

```

# zfs send -I pool/fs@snap1 pool/clone@snapA > /snaps/fscloonesnap-I
.
.
# zfs receive -F pool/clone < /snaps/fscloonesnap-I

```

- You can send a replication stream of all descendent file systems, up to the named snapshots. When received, all properties, snapshots, descendent file systems, and clones are preserved. For example:

```
# zfs send -R pool/fs@snap > snaps/fs-R
```

- Send an incremental replication stream.

```
zfs send -R -[iI] @snapA pool/fs@snapD
```

For extended examples, see the *Solaris ZFS Administration Guide*.

- **ZFS quotas and reservations for file system data only** – In the Solaris 10 10/08 release, dataset quotas and reservations are provided that do not include descendents, such as snapshots and clones, in the space consumption accounting. The existing ZFS quota and reservation features remain as in previous Solaris releases.
 - The `refquota` property limits the amount of space a dataset can consume. This property enforces a hard limit on the amount of space that can be used. This hard limit does not include space used by descendents, such as snapshots and clones.
 - The `refreservation` property sets the minimum amount of space that is guaranteed to a dataset, not including its descendents.

For example, you can set a 10-Gbyte `refquota` for `studentA` that sets a 10-Gbyte hard limit of *referenced* space. For additional flexibility, you can set a 20-Gbyte quota that allows you to manage `studentA`'s snapshots.

```

# zfs set refquota=10g tank/studentA
# zfs set quota=20g tank/studentA

```

- **ZFS storage pool properties** – New ZFS storage pool property information is provided in the Solaris 10 10/08 release.

- Display all pool attributes – You can use the `zpool get all pool` command to display all pool property information. For example:

```
# zpool get all users
NAME  PROPERTY  VALUE      SOURCE
users size      16.8G      -
users used    194K       -
users available 16.7G      -
users capacity 0%         -
users altroot -           default
users health  ONLINE     -
users guid    14526624140147884971 -
users version 10         default
users bootfs  -          default
users delegation on         default
users autoreplace off        default
users cachefile -          default
users failmode wait       default
```

- The `cachefile` property – This release provides the `cachefile` property, which controls where pool configuration information is cached. All pools in the cache are automatically imported when the system boots. However, installation and clustering environments might need to cache this information in a different location so that pools are not automatically imported.

You can set this property to cache pool configuration in a different location that can be imported later by using the `zpool import -c` command. For most ZFS configurations, this property would not be used.

The `cachefile` property is not persistent and is not stored on disk. This property replaces the temporary property that was used to indicate that pool information should not be cached in previous Solaris releases.

- The `failmode` property – This release provides the `failmode` property for determining the behavior of a catastrophic pool failure due to a loss of device connectivity or the failure of all devices in the pool. The `failmode` property can be set to these values: `wait`, `continue`, or `panic`. The default value is `wait`, which means you must reconnect the device or replace a failed device and clear the error with the `zpool clear` command.

The `failmode` property is set like other settable ZFS properties, which can be set either before or after the pool is created. For example:

```
# zpool set failmode=continue tank
# zpool get failmode tank
NAME  PROPERTY  VALUE      SOURCE
tank  failmode  continue    local

# zpool create -o failmode=continue users mirror c0t1d0 c1t1d0
```

- **ZFS and file system mirror mounts** – In the Solaris 10 10/08 release, NFSv4 mount enhancements are provided to make ZFS file systems more accessible to NFS clients. When file systems are created on the NFS server, the NFS client can automatically discover these newly created file systems within their existing mount of a parent file system. For example, if the server neo already shares the tank file system and client zee has it mounted, /tank/baz is automatically visible on the client after it is created on the server.

```
zee# mount neo:/tank /mnt
zee# ls /mnt
baa  bar
```

```
neo# zfs create tank/baz
```

```
zee% ls /mnt
baa  bar  baz
zee% ls /mnt/baz
file1  file2
```

- **ZFS command history enhancements** (zpool history) – In the Solaris 10 10/08 release, the zpool history command provides the following new features:

- ZFS file system event information is displayed. For example:

```
# zpool history users
History for 'users':
2008-07-10.09:43:05 zpool create users mirror c1t1d0 c1t2d0
2008-07-10.09:43:48 zfs create users/home
2008-07-10.09:43:56 zfs create users/home/markm
2008-07-10.09:44:02 zfs create users/home/marks
2008-07-10.09:44:19 zfs snapshot -r users/home@yesterday
```

- A -l option for displaying a long format that includes the user name, the host name, and the zone in which the operation was performed. For example:

```
# zpool history -l users
History for 'users':
2008-07-10.09:43:05 zpool create users mirror c1t1d0 c1t2d0
[user root on corona:global]
2008-07-10.09:43:13 zfs create users/marks
[user root on corona:global]
2008-07-10.09:43:44 zfs destroy users/marks
[user root on corona:global]
2008-07-10.09:43:48 zfs create users/home
[user root on corona:global]
2008-07-10.09:43:56 zfs create users/home/markm
[user root on corona:global]
2008-07-10.09:44:02 zfs create users/home/marks
[user root on corona:global]
```

```
2008-07-11.10:44:19 zfs snapshot -r users/home@yesterday
[user root on corona:global]
```

- A `-i` option for displaying internal event information that can be used for diagnostic purposes. For example:

```
# zpool history -i users
History for 'users':
2008-07-10.09:43:05 zpool create users mirror c1t1d0 c1t2d0
2008-07-10.09:43:13 [internal create txg:6] dataset = 21
2008-07-10.09:43:13 zfs create users/marks
2008-07-10.09:43:48 [internal create txg:12] dataset = 27
2008-07-10.09:43:48 zfs create users/home
2008-07-10.09:43:55 [internal create txg:14] dataset = 33
2008-07-10.09:43:56 zfs create users/home/markm
2008-07-10.09:44:02 [internal create txg:16] dataset = 39
2008-07-10.09:44:02 zfs create users/home/marks
2008-07-10.09:44:19 [internal snapshot txg:21] dataset = 42
2008-07-10.09:44:19 [internal snapshot txg:21] dataset = 44
2008-07-10.09:44:19 [internal snapshot txg:21] dataset = 46
2008-07-10.09:44:19 zfs snapshot -r users/home@yesterday
```

- **Upgrading ZFS file systems** (`zfs upgrade`) – In the Solaris 10 10/08 release, you can use the `zfs upgrade` command to upgrade your existing ZFS file systems with new file system enhancements. ZFS storage pools have a similar upgrade feature to provide pool enhancements to existing storage pools.

For example:

```
# zfs upgrade
This system is currently running ZFS filesystem version 2.
```

The following filesystems are out of date, and can be upgraded. After being upgraded, these filesystems (and any 'zfs send' streams generated from subsequent snapshots) will no longer be accessible by older software versions.

```
VER  FILESYSTEM
---  -----
1   datab
1   datab/users
1   datab/users/area51
```

Note – File systems that are upgraded and any streams that are created from those upgraded file systems by the `zfs send` command are not accessible on systems that are running older software releases.

- **ZFS delegated administration** – In the Solaris 10 10/08 release, you can delegate granular permissions to perform ZFS administration tasks to non-privileged users.

You can use the `zfs allow` and `zfs unallow` commands to grant and remove permissions. You can modify the ability to use delegated administration with the pool's `delegation` property. For example:

```
# zpool get delegation users
NAME PROPERTY  VALUE      SOURCE
users delegation on         default
# zpool set delegation=off users
# zpool get delegation users
NAME PROPERTY  VALUE      SOURCE
users delegation off         local
```

By default, the `delegation` property is enabled.

- **Setting up separate ZFS logging devices** – The ZFS intent log (ZIL) is provided to satisfy POSIX requirements for synchronous transactions. For example, databases often require their transactions to be on stable storage devices when returning from a system call. NFS and other applications can also use `fsync()` to ensure data stability. By default, the ZIL is allocated from blocks within the main storage pool. However, better performance in the Solaris 10 10/08 release might be possible by using separate ZIL devices in your ZFS storage pool, such as with NVRAM or a dedicated disk,

Log devices for the ZIL are not related to database log files.

You can set up a ZFS logging device when the storage pool is created or after the pool is created. For examples of setting up log devices, see *Solaris ZFS Administration Guide*.

- **Creating intermediate ZFS datasets** – In the Solaris 10 10/08 release, you can use the `-p` option with the `zfs create`, `zfs clone`, and `zfs rename` commands to quickly create a nonexistent intermediate dataset, if it doesn't already exist.

For example, create ZFS datasets (`users/area51`) in the `datab` storage pool.

```
# zfs list
NAME                                USED AVAIL REFER MOUNTPOINT
datab                                106K 16.5G 18K /datab
# zfs create -p -o compression=on datab/users/area51
```

If the intermediate dataset exists during the create operation, the operation completes successfully.

Properties specified apply to the target dataset, not to the intermediate datasets. For example:

```
# zfs get mountpoint,compression datab/users/area51
NAME                                PROPERTY  VALUE      SOURCE
datab/users/area51 mountpoint /datab/users/area51 default
datab/users/area51 compression on         local
```

The intermediate dataset is created with the default mount point. Any additional properties are disabled for the intermediate dataset. For example:

```
# zfs get mountpoint,compression datab/users
NAME          PROPERTY      VALUE          SOURCE
datab/users   mountpoint    /datab/users   default
datab/users   compression   off            default
```

For more information, see [zfs\(1M\)](#).

- **ZFS hot-plugging enhancements** – In the Solaris 10 10/08 release, ZFS more effectively responds to removed devices and provides a mechanism to automatically identify devices that are inserted:
 - You can replace an existing device with an equivalent device without having to use the `zpool replace` command.

The `autoreplace` property controls automatic device replacement. If set to `off`, device replacement must be initiated by the administrator by using the `zpool replace` command. If set to `on`, any new device that is found in the same physical location as a device that previously belonged to the pool, is automatically formatted and replaced. The default behavior is `off`.
 - The storage pool state `REMOVED` is provided when a device or hot spare has been removed, if the device was physically removed while the system was running. A hot spare device is substituted for the removed device, if available.
 - If a device is removed and then inserted, the device is placed online. If a hot spare was activated when the device is re-inserted, the hot spare is removed when the online operation completes.
 - Automatic detection when devices are removed or inserted is hardware-dependent and might not be supported on all platforms. For example, USB devices are automatically configured upon insertion. However, you might have to use the `cfgadm -c configure` command to configure a SATA drive.
 - Hot spares are checked periodically to make sure that they are online and available.

For more information, see [zpool\(1M\)](#).

- **Recursively renaming ZFS snapshots** (`zfs rename -r`) – In the Solaris 10 10/08 release, you can recursively rename all descendent ZFS snapshots by using the `zfs rename -r` command. For example, snapshot a set of ZFS file systems.

```
# zfs snapshot -r users/home@today
# zfs list
NAME                                USED AVAIL REFER MOUNTPOINT
users                                216K 16.5G 20K  /users
users/home                          76K 16.5G 22K  /users/home
users/home@today                    0    -    22K  -
users/home/markm                    18K 16.5G 18K  /users/home/markm
```

```

users/home/markm@today      0      -    18K  -
users/home/marks           18K   16.5G  18K  /users/home/marks
users/home/marks@today     0      -    18K  -
users/home/neil            18K   16.5G  18K  /users/home/neil
users/home/neil@today      0      -    18K  -

```

Then, rename the snapshots the following day.

```

# zfs rename -r users/home@today @yesterday
# zfs list
NAME                                USED  AVAIL  REFER  MOUNTPOINT
users                                216K  16.5G  20K    /users
users/home                           76K   16.5G  22K    /users/home
users/home@yesterday                 0      -    22K    -
users/home/markm                     18K   16.5G  18K    /users/home/markm
users/home/markm@yesterday           0      -    18K    -
users/home/marks                     18K   16.5G  18K    /users/home/marks
users/home/marks@yesterday           0      -    18K    -
users/home/neil                      18K   16.5G  18K    /users/home/neil
users/home/neil@yesterday            0      -    18K    -

```

Snapshots are the only datasets that can be renamed recursively.

- **GZIP compression now available for ZFS** – In the Solaris 10 10/08 release, you can set gzip compression on ZFS file systems in addition to lzjb compression. You can specify compression as gzip, the default, or gzip-*N*, where *N* equals 1 through 9. For example:

```

# zfs create -o compression=gzip users/home/snapshots
# zfs get compression users/home/snapshots
NAME                                PROPERTY  VALUE      SOURCE
users/home/snapshots                compression  gzip       local
# zfs create -o compression=gzip-9 users/home/oldfiles
# zfs get compression users/home/oldfiles
NAME                                PROPERTY  VALUE      SOURCE
users/home/oldfiles                  compression  gzip-9     local

```

- **Storing multiple copies of ZFS user data** – A ZFS file system automatically stores metadata multiple times across different disks, if possible, as a reliability feature. This feature is known as *ditto blocks*. In the Solaris 10 10/08 release, you can specify that multiple copies of user data is also stored per file system by using the `zfs set copies` command. For example:

```

# zfs set copies=2 users/home
# zfs get copies users/home
NAME      PROPERTY  VALUE  SOURCE
users/home  copies    2      local

```

Available values are 1, 2, or 3. The default value is 1. These copies are in addition to any pool-level redundancy, such as in a mirrored or RAID-Z configuration.

For more information about using this property, see the [Solaris ZFS Administration Guide](#).

- **ZFS command history** (zpool history) – In the Solaris 10 8/07 release, ZFS automatically logs successful zfs and zpool commands that modify pool state information. This feature enables you or Sun support personnel to identify the *exact* ZFS commands that were executed to troubleshoot an error scenario.
- **Improved storage pool status information** (zpool status) – In the Solaris 10 8/07 release, you can use the `zpool status -v` command to display a list of files with persistent errors. Previously, you had to use the `find -inum` command to identify the file names from the list of displayed inodes.
- **ZFS and Solaris iSCSI improvements** – In the Solaris 10 8/07 release, you can create a ZFS volume as a Solaris iSCSI target device by setting the `shareiscsi` property on the ZFS volume. This method is a convenient way to quickly set up a Solaris iSCSI target. For example:

```
# zfs create -V 2g tank/volumes/v2
# zfs set shareiscsi=on tank/volumes/v2
# iscsitadm list target
Target: tank/volumes/v2
       iSCSI Name: iqn.1986-03.com.sun:02:984fe301-c412-ccc1-cc80-cf9a72aa062a
       Connections: 0
```

After the iSCSI target is created, you would set up the iSCSI initiator. For information about setting up a Solaris iSCSI initiator, see [Chapter 14, “Configuring Solaris iSCSI Targets and Initiators \(Tasks\)”](#) in *System Administration Guide: Devices and File Systems*.

For more information about managing a ZFS volume as an iSCSI target, see the *Solaris ZFS Administration Guide*.

- **ZFS property improvements**
 - **ZFS `xattr` property** – In the Solaris 10 8/07 release, you can use the `xattr` property to disable or enable extended attributes for a specific ZFS file system. The default value is `on`.
 - **ZFS `canmount` property** – In the Solaris 10 8/07 release, you use the `canmount` property to specify whether a dataset can be mounted by using the `zfs mount` command.
 - **ZFS user properties** – In the Solaris 10 8/07 release, ZFS supports user properties, in addition to the standard native properties that can either export internal statistics or control ZFS file system behavior. User properties have no effect on ZFS behavior, but you can use them to annotate datasets with information that is meaningful in your environment.
 - **Setting properties when creating ZFS file systems** – In the Solaris 10 8/07 release, you can set properties when you create a file system, in addition to setting properties after the file system is created.

The following examples illustrate equivalent syntax:

```
# zfs create tank/home
# zfs set mountpoint=/export/zfs tank/home
# zfs set sharenfs=on tank/home
# zfs set compression=on tank/home
```

Or, set the properties when the file system is created.

```
# zfs create -o mountpoint=/export/zfs -o sharenfs=on -o
compression=on tank/home
```

- **Display all ZFS file system information** – In the Solaris 10 8/07 release, you can use various forms of the `zfs get` command to display information about all datasets if you do not specify a dataset. In previous releases, all dataset information was not retrievable with the `zfs get` command.

For example:

```
# zfs get -s local all
tank/home          atime          off            local
tank/home/bonwick atime          off            local
tank/home/marks    quota         50G           local
```

- **New `zfs receive -F` option** – In the Solaris 10 8/07 release, you can use the new `-F` option to the `zfs receive` command to force a rollback of the file system to the most recent snapshot before doing the receive operation. Using this option might be necessary when the file system is modified between the time a rollback occurs and the receive operation is initiated.
- **Recursive ZFS snapshots** – In the Solaris 10 11/06 release, recursive snapshots are available. When you use the `zfs snapshot` command to create a file system snapshot, you can use the `-r` option to recursively create snapshots for all descendant file systems. In addition, using the `-r` option recursively destroys all descendant snapshots when a snapshot is destroyed.
- **Double Parity RAID-Z (raidz2)** – In the Solaris 10 11/06 release, replicated RAID-Z configuration can now have either single-parity or double-parity, which means that one or two device failures can be sustained respectively, without any data loss. You can specify the `raidz2` keyword for a double-parity RAID-Z configuration. Or, you can specify the `raidz` or `raidz1` keyword for a single-parity RAID-Z configuration.
- **Hot spares for ZFS storage pool devices** – Starting in the Solaris 10 11/06 release, the ZFS hot spares feature enables you to identify disks that could be used to replace a failed or faulted device in one or more storage pools. Designating a device as a *hot spare* means that if an active device in the pool fails, the hot spare automatically replaces the failed device. Or, you can manually replace a device in a storage pool with a hot spare.
- **Replacing a ZFS file system with a ZFS clone (`zfs promote`)** – In the Solaris 10 11/06 release, the `zfs promote` command enables you to replace an existing ZFS file system with a clone of that file system. This feature is helpful when you want to run tests on an alternative version of a file system and then, make that alternative version of the file system the active file system.

- **ZFS and zones improvements** – In the Solaris 10 11/06 release, the ZFS and zones interaction is improved. On a Solaris system with zones installed, you can use the `zoneadm clone` feature to copy the data from an existing source ZFS zonepath to a target ZFS zonepath on your system. You cannot use the ZFS clone feature to clone the non-global zone. You must use the `zoneadm clone` command. For more information, see *System Administration Guide: Solaris Containers-Resource Management and Solaris Zones*.
- **Upgrading ZFS storage pools** (`zpool upgrade`) – Starting in the Solaris 10 6/06 release, you can upgrade your storage pools to a newer version to take advantage of the latest features by using the `zpool upgrade` command. In addition, the `zpool status` command has been modified to notify you when your pools are running older versions.
- **Clearing device errors** – Starting in the Solaris 10 6/06 release, you can use the `zpool clear` command to clear error counts that are associated with a device or the pool. Previously, error counts were cleared when a device in a pool was brought online with the `zpool online` command.
- **Recovering destroyed pools** – In the Solaris 10 6/06 release, the `zpool import -D` command enables you to recover pools that were previously destroyed with the `zpool destroy` command.
- **ZFS backup and restore commands renamed** – In the Solaris 10 6/06 release, the `zfs backup` and `zfs restore` commands are renamed to `zfs send` and `zfs receive` to more accurately describe their function. Their function is to save and restore ZFS data stream representations.
- **Compact NFSv4 ACL format** – Starting in the Solaris 10 6/06 release, three NFSv4 ACL formats are available: verbose, positional, and compact. The new compact and positional ACL formats are available to set and display ACLs. You can use the `chmod` command to set all three ACL formats. Use the `ls -V` command to display compact and positional ACL formats. Use the `ls -v` command to display verbose ACL formats.
- **Temporarily take a device offline** – Starting in the Solaris 10 6/06 release, you can use the `zpool offline -t` command to take a device offline temporarily. When the system is rebooted, the device is automatically returned to the ONLINE state.
- **ZFS is integrated with Fault Manager** – Starting in the Solaris 10 6/06 release, a ZFS diagnostic engine that is capable of diagnosing and reporting pool failures and device failures is included. Checksum, I/O, and device errors associated with pool or device failures are also reported. Diagnostic error information is written to the console and the `/var/adm/messages` file. In addition, detailed information about recovering from a reported error can be displayed by using the `zpool status` command.

For more information about these improvements and changes, see the *Solaris ZFS Administration Guide*.

See the following What's New sections for related ZFS feature information:

- “The Solaris ZFS File System” on page 123
- “File-System Monitoring Tool” on page 110

- “Improved Device in Use Error Checking” on page 116

nss_ldap shadowAccount **Support**

The LDAP name service is enhanced to support the account locking and password aging functionality using the data in the shadow database stored on a configured LDAP server. This support enables the `passwd(1)` utility and the `pam_unix_*(5)` PAM modules to function almost the same when handling account locking and password aging for local accounts and remote LDAP user accounts. Therefore, using the `pam_ldap(5)` module is no longer the only way to implement the password policy and account control for the LDAP name service. `pam_unix_*(5)` can be used to obtain the same consistent results as with the files and `nisplus` name services.

For more information, see *System Administration Guide: Naming and Directory Services (DNS, NIS, and LDAP)*.

Sun Validation Test Suite 7.0 Patch Set 6

The SunVTS™ 7.0 Patch Set 6 is integrated in the Solaris 10 10/09 release. The SunVTS 7.0 Patch Set 6 follows a conventional three-tier architecture model. The patch set includes a browser-based user interface (BUI), a Java technology-based middle server, and a diagnostic agent. Enhancements to the SunVTS infrastructure include:

- Support for solid-state drives (SSD) added to `vtsk`
- Default level of logical tests enhanced to adapt to system configuration size
- Minimum and maximum values or hard limit for reserve swap in `vtsk`
- Ability to change the sequence of Logical Test execution

The Solaris 10 10/09 release includes the following enhancements to memory and CPU diagnostics:

- Coverage is added for X86-L3\$ in `l3sramtest`
- Enhanced `vmemtest`, `fputest`, and `l2sramtest` provide callbacks to return swap requirements
- Tuned logical tests for x86 systems and UltraSPARC® T2 Processor-based systems

The Solaris 10 10/09 release also includes the following enhancements to the I/O diagnostics:

- `disktest` is enhanced to run in Read only mode if Write or Read option is not applicable
- Disk Logical Test is tuned for x86, UltraSPARC T2 Processor, and UltraSPARC IV systems
- `disktest` options are automated to run solid-state drive (SSD) and hard disk drive (HDD) Tasks in Disk LT
- Selection of test-options is automated in `netlbttest`
- Support in `disktest` and `iobustest` for safe and unsafe test options

Installation Enhancements

The following installation feature has been added to the Solaris 10 10/09 release.

Turbo-Charging SVR4 Packaging

Starting with the Solaris 10 10/09 release, SVR4 package commands run faster. This enhancement means that the Solaris installation technologies, such as initial installations, upgrades, live upgrades, and zone installations, perform significantly faster.

System Resources Enhancements

The following system resources feature has been added to the Solaris 10 10/09 release.

Zones Parallel Patching

The zones parallel patching enhancement to the standard Solaris 10 patch utilities increases the patching tools performance on systems with multiple zones by allowing parallel patching of the non-global zones. In releases prior to Solaris 10 10/09, the feature is implemented through the patch utilities patch 119254-66 or later revision for SPARC and 119255-66 or later revision for x86. The global zone is still patched before the non-global zones are patched.

For more information, see the following:

- <http://blogs.sun.com/patch/date/20090619>
- *System Administration Guide: Solaris Containers-Resource Management and Solaris Zones*

PVIO Drivers in Solaris 10 Guest Domains

When using the Sun xVM hypervisor in a Solaris OS, fully virtualized guest domains are referred to as hardware-assisted virtual machines (HVMs). HVM + PVIO guests provide better performance through the use of PV drivers.

Releases starting with the Solaris 10 10/08 release are shipped with the Solaris PV drivers. A patch is available for Solaris 10 5/08.

For more information, see “Solaris 10 releases” in “Guests That Are Known to Work” in *System Administration Guide: Virtualization Using the Solaris Operating System*. This guide also discusses HVM-capable machines.

Device Management Enhancements

The following device management feature has been added to the Solaris 10 10/09 release.

iSCSI Initiator SMF Service

A new SMF service, under FMRI `svc:/network/iscsi/initiator:default` is introduced to control the availability of iSCSI devices. The SMF service also controls the timing to start discovery and enumeration of iSCSI devices during OS start—up.

Other services that rely on the availability of iSCSI devices can customize their dependency on the this new iSCSI Initiator services. For more information see, the `iscsi(7D)` man page.

LSI 6180 Controller Support in Solaris MPxIO

Starting with the Solaris 10 10/09 release, Solaris MPxIO supports the LSI 6180 controller-based storage arrays.

System Performance Enhancements

The following system performance feature has been added to the Solaris 10 10/09 release.

Callout Subsystem Scalability

The callout subsystem is redesigned to include the following features:

- Performance and scalability improvements:
 - Per-CPU data structures to minimize mutex contention
 - Per-CPU callout processing to improve scalability
 - Event-based implementation that avoids polling overhead
- High-resolution timers for improved functionality. Many API calls use high-resolution timers and do not experience latency because the system rounds off the specified intervals. These timers include commonly used calls such as `poll()` and `nanosleep()`.
- Observability improvements:
 - Comprehensive set of options for the `MDB dcmd callout`
 - New `MDB dcmd callout id`
 - New callout `kstats`

Driver Enhancements

The following driver features and enhancements have been added to the Solaris 10 10/09 release.

Enhanced 10 Gigabit Ethernet Performance

The Solaris 10 5/09 includes many enhancements to the Solaris 10GbE drivers. The `nxge` 10GbE driver includes the following enhancements:

- TCP receive throughput is improved from 40% at 8 connections to over 90% improvements for 32, 100, 400, and 1000 connections
- TCP transmit throughput is improved from almost 80% at 8 connections to over 100% for the higher connection tests
- UDP transmit throughput is improved from 80% for 64 byte messages to over 160% for 8 Kbyte messages

The `ixgbe` driver on x86 systems includes the following enhancements:

- TCP transmit throughput is improved to almost 100% for 8 or more connections
- TCP receive rates are 10Gb line rate for 8, 32, 100, 400, and 1000 connections

- UDP transmit maximum throughput doubles to 10Gb line rate
- Ping pong data rates are improved from 2x to 3x as the message size increases from 64 bytes to 512 bytes

Solaris 10GbE drivers can now deliver close to line data rates providing optimal performance on 10 Gigabit networks.

InfiniBand Feature Improvements

The Solaris 10 5/09 release includes the following InfiniBand-related enhancements:

- **InfiniBand Host Channel Adapter (HCA)** – The Solaris 10 5/09 release includes a significantly enhanced InfiniBand driver for the Mellanox ConnectX HCA. The InfiniBand driver enables InfiniBand protocols to operate over both Double Data Rate (DDR) and Quad Data Rate (QDR) InfiniBand fabrics. The driver is also integrated into the Solaris FMA framework for fault management and the driver supports relaxed ordering on SPARC systems.
- **InfiniBand Transport Framework (IBTF)** – The Solaris 10 5/09 release includes a significantly improved IBTF implementation that provides enhanced support for running RDMA-based InfiniBand protocols in Solaris. InfiniBand for SPARC now supports both the PCI Dynamic Reconfiguration (DR).
- **Internet Protocol over InfiniBand (IPoIB)** – The Solaris 10 5/09 release includes a significantly improved IPoIB driver (`ibd`) supporting Internet RFCs 4391 and 4392. The IPoIB driver in the Solaris 10 5/09 release supports the User Datagram (UD) mode of operation, IPv4 and IPv6 addressing, and takes advantage of hardware offloads in the ConnectX HCA for improved throughput at lower CPU utilization. IPoIB-UD enables the use of any TCP/IP application protocol such as SSH, HTTP, FTP, NFS, and iSCSI over both Double Data Rate (DDR) and Quad Data Rate (QDR) InfiniBand fabrics. The new IPoIB driver for both SPARC and x86 platforms offers a significant performance boost over the previously available driver.
- **Sockets Direct Protocol (SDP)** – The Solaris 10 5/09 release includes a significantly improved SDP driver and `sockfs` implementation. SDP is a transport protocol layered over the Infiniband Transport Framework (IBTF). SDP is a standard implementation based on Annex 4 of the Infiniband Architecture Specification Vol 1. The SDP protocol provides reliable byte-stream, flow controlled two-way data transmission that is similar to the Transmission Control Protocol (TCP). InfiniBand programmers use SDP through the `libsdp` C library that supports a sockets-based `SOCK_STREAM` interface to application programs. The SDP protocol supports graceful close, IPv4 and IPv6 addressing, the connecting/accepting connect model, out-of-band (OOB) data and common socket options. The SDP protocol also supports kernel bypass data transfers and data transfers from send-upper-layer-protocol (ULP) buffers to receive ULP buffers.
- **Reliable Datagram Sockets (RDS)** – The Solaris 10 5/09 release includes an improved RDSv1 driver certified for use with Oracle RAC (Real Application Clusters) 10gR2.

- **User-Level Direct Access Programming Library (uDAPL)** – The Solaris 10 5/09 release includes an updated uDAPL over InfiniBand API that conforms to the latest Direct Access Transport (DAT) Collaborative uDAPL 1.2 specification.

Driver for LSI MPT 2.0-Compliant SAS2.0 Controllers

The `mpt_sas(7D)` driver supports SAS, SATA, SMP physical devices, and virtual devices by using the Integrated RAID feature. The new architecture for SAS drivers supports the following features:

- SAS initiator ports (iports)
- Dynamic reconfiguration of SAS, SATA, and SMP targets
- FWARC 2008/013-compliant device representation
- Multipathing

For more information, see the `mpt_sas(7D)` man page.

x86: Broadcom NetXtreme II Gigabit Ethernet Driver

The Solaris 10 10/09 release includes new chipsets support such as `bcm5716c` and `bcm5716s`.

x86: Interrupt Remapping Support for Intel Vt-d

The Solaris 10 10/09 release provides an interrupt-remapping table that isolates interrupts on at least the Intel Nehalem platform and ensures that devices can only use authorized interrupts and that the interrupts are properly targeted. This feature improves system reliability, availability, and serviceability (RAS).

x86: SATA Tape Devices Supported by the AHCI Driver

SATA tape devices are now supported by the AHCI driver. Users can connect or hot-plug the SATA tape drive to the AHCI controller through the SATA or eSATA cable. Error handling mechanism is also enhanced for SATA ATAPI devices including CD, DVD, or tape.

For more information, see the `ahci(7D)` man page.

Sun StorageTek 6Gb/s SAS PCIe RAID HBA Driver

The `mr_sas` MegaRAID SAS2.0 controller host bus adapter driver is a SCSI-compliant nexus driver that supports the LSI MegaRAID SAS 92xx series of controllers and the StorageTek 6Gb/s SAS RAID HBA series of controllers and the LSI MegaRAID SAS 92xx series of controllers.

Some of the supported RAID features include:

- RAID levels 0, 1, 5, and 6, and RAID spans 10, 50 and 60
- Online capacity expansion (OCE)
- Online RAID level migration (RLM)

- Auto resume after loss of system power during array rebuild or reconstruction (OCE or RLM)
- Configurable stripe size up to 1 Mbyte
- Capability to check consistency for background data integrity
- Patrol read for media scanning and repairing
- 64 logical drive support
- Up to 64 TB LUN support
- Automatic rebuild, and global and dedicated hot—spare support

Intel 82599 10Gb PCI Express Ethernet Controller

Starting with the Solaris 10 10/09 release, the ixgbe driver supports the Intel 82599 10Gb PCI Express Ethernet Controller chipset.

Intel 82598 10Gb PCI Express Ethernet Controller

Starting with the Solaris 10 10/09 release, the ixgbe driver supports the Intel 82598 10Gb PCI Express Ethernet Controller chipset.

Freeware Enhancements

The following freeware features and enhancements have been added to the Solaris 10 10/09 release.

NTP Version 4.2.5

The Solaris 10 10/09 release includes the latest version of the network time protocol that supports enhanced authentication, IPv6, and greater performance. For more information, see the [ntptime\(1M\)](#) man page.

PostgreSQL

The Solaris 10 10/09 release supports PostgreSQL versions 8.1.17, 8.2.13, and 8.3.7.

Samba

The Solaris 10 10/09 release supports Samba 3.0.35.

What's New in the Solaris 10 5/09 Release

This chapter summarizes new features in the Solaris 10 5/09 release.

System Resources Enhancements

The following system resources features and enhancements have been added to the Solaris 10 5/09 release.

Support Added for Using ZFS Clones When Cloning a Zone

If the source and the target zonpaths reside on ZFS and both are in the same pool, a snapshot of the source zonpath is taken and the zoneadm clone uses ZFS to clone the zone.

You can specify to copy a ZFS zonpath instead of specifying to clone the ZFS. If neither the source nor the target zonpath is on ZFS, or if one is on ZFS and the other is not on ZFS, the clone process uses the existing copy technique.

In all cases, the system copies the data from a source zonpath to a target zonpath if using a ZFS clone is not possible.

For more information, see the following:

- [zoneadm\(1M\)](#) man page
- *System Administration Guide: Solaris Containers-Resource Management and Solaris Zones*

zoneadm attach -b Option

Use the -b option to specify official or Interim Diagnostics Relief (IDR) patches, to be backed out of a zone during the attach. This option applies only to zone brands that use SVr4 packaging.

For more information, see the following:

- [zoneadm\(1M\)](#) man page
- *System Administration Guide: Solaris Containers-Resource Management and Solaris Zones*

System Administration Enhancements

The following system administration features and enhancements have been added to the Solaris 10 5/09 release.

SMF Services for IPsec

IP security (IPsec) is now managed by the following Solaris Management Facility (SMF) services:

- `svc:/network/ipsec/policy:default` – The `policy` service checks for the `/etc/inet/ipsecinit.conf` file and feeds the data into the IPsec Security Policy Database (SPD). The `policy` service must be started and its file, `/etc/inet/ipsecinit.conf`, must exist for boot-time IPsec policy configuration.
- `svc:/network/ipsec/ike:default` – The `ike` service controls the Internet Key Exchange (IKE) daemon in `iked(1M)`. This service controls `ike` in a manner similar to other daemon-controlled services like `ssh` or `sendmail`.
- `svc:/network/ipsec/manual-key:default` – The `manual-key` service checks for the `/etc/inet/secret/ipseckeys` file and feeds the keys into the IPsec Security Association Database (SADB). Prior to SMF, the mere existence of the `/etc/inet/secret/ipseckeys` file was sufficient, but now the service should also be enabled to load manual IPsec keys.
- `svc:/network/ipsec/ipsecalgs:default` – The `ipsecalgs` service is enabled by default and maps Solaris Cryptographic Framework algorithms to their use in IPsec. Changes enabled with `ipsecalgs(1M)` subsequently refresh the `ipsecalgs` service.

The SMF management brings all the SMF features to IPsec, for example, interface consistency, capability of restarting, and fault-tracking.

Security Enhancements

The following security features and enhancements have been added to the Solaris 10 5/09 release.

NAT-Traversal for IPsec Key Management Developers

The Solaris 10 5/09 release contains a public API for User Datagram Protocol (UDP) sockets that act as IPsec Network Address Translator (NAT) Traversal endpoints.

The `UDP_NAT_T_ENDPOINT` socket option, when enabled, has UDP traffic prefixed with a zero security parameters index (SPI) value of four bytes on outbound traffic and strips zero SPIs on inbound traffic. Inbound traffic bound for such a socket with a nonzero SPI is automatically transferred to IPsec's Encapsulating Security Payload (ESP) for ESP-in-UDP decapsulation. ESP-in-UDP encapsulation is determined by a property in the IPsec Security Association (SA).

This feature enables IPsec key management software developers to create key management protocols that can transit NAT devices. The Solaris IKE daemon in `iked(1M)` uses this facility and such sockets are displayed using the `pfules(1M)` command.

Stronger Algorithms for IPsec

The Solaris 10 5/09 release introduces the following algorithms for IPsec and IKE:

- Three larger Diffie-Hellman integer-modulus groups including 2048-bit, 3072-bit, and 4096-bit – The larger Diffie-Hellman groups are available in IKE Phase 1 and Phase 2. The groups are specified by group number 14 for 2048-bit, 15 for 3072-bit, and 16 for 4096-bit, per RFC 3526.
- SHA-2 series of hashes including sha256, sha384, and sha512– SHA-2 using HMAC is available for IPsec's Authentication Header (AH) and ESP, and for IKE during its interactions. SHA-2 is used in IPsec per RFC 4868, with truncated ICV lengths of 16 bytes for SHA256, 24 bytes for SHA384, and 32 bytes for SHA512.

Note – SHA-2 is not available for certificates generated with `ikecert(1M)`.

SunSSH With OpenSSL PKCS#11 Engine Support

This feature enables the SunSSH server and client to use Solaris Cryptographic Framework through the OpenSSL PKCS#11 engine. SunSSH uses cryptographic framework for hardware crypto acceleration of symmetric crypto algorithms which is important to the data transfer speed. This feature is aimed at UltraSPARC® T2 processor platforms with `n2cp(7D)` crypto driver.

UltraSPARC T1 processor platforms are not affected by this feature since the `ncp(7D)` driver does not support symmetric crypto algorithms. Platforms without any hardware crypto plugins are not affected by this feature, regardless of the value set for the `UseOpenSSLEngine` option. The default value of the `UseOpenSSLEngine` option is set to `on` and the server and client SSH configuration files need not be updated.

SunSSH should be used with Sun Crypto Accelerator 6000 board software version 1.1 with the following patches installed:

- 128365-02 for SPARC-based systems
- 128366-02 for x86-based systems

Note – No patch is available for the Sun Crypto Accelerator 6000 board software version 1.0. To workaround this issue, remove the AES counter modes from the `Ciphers` option keyword on both the server and the client side.

For more information, see the [ssh_config\(4\)](#) and [sshd_config\(4\)](#)

Device Management Enhancements

The following device management feature has been added to the Solaris 10 5/09 release.

x86: T-State Support for Intel Based Processors

This feature provides the basic CPU Advanced Configuration and Power Interface (ACPI) T-state support. T-state support enables the CPU driver to receive `_TPC` change notifications as a manner of controlling the processor speed. This is frequently done on some systems as a passive cooling mechanism along with the existing CPU ACPI P-States.

For more information, see <http://opensolaris.org/os/community/pm/>.

System Performance Enhancements

The following system performance features and enhancements have been added to the Solaris 10 5/09 release.

Large Segment Offload Support for Intel PCI Express 10Gb NIC Driver

This feature introduces Large Segment Offload (LSO) support for the `ixgbe` driver and some `ixgbe` driver bug fixes. LSO is an important feature for NIC, especially for 10-Gb NIC. LSO can offload the segmentation job on Layer 4 to the NIC driver. LSO improves transmit performance by decreasing CPU overhead. This feature is enabled by default.

Solaris Power Aware Dispatcher and Deep C-State Support

This feature includes the following enhancements:

- Event driven CPU power management – On systems that support Dynamic Voltage and Frequency Scaling (DVFS) by Solaris, the kernel scheduler or dispatcher will schedule threads across the system's CPUs in a manner that coalesces load, and frees up other CPUs to be deeply power managed. CPU power state changes are triggered when the dispatcher recognizes that the utilization across a group of power manageable CPUs has changed in a significant way. This eliminates the need to periodically poll CPU utilizations across the system, and enables the system to save more power when CPUs are not used, while driving performance when CPUs are used. Event driven CPU power management is enabled by default on systems that support DVFS. This feature can be disabled, or the legacy polling-based CPU power management can be used through the `cpupm` keyword in `power.conf(4)`.
- Support for Deep Idle CPU Power Management or deep C-state support on Intel Nehalem-based systems – The project also adds Solaris support for Deep C-states on Intel Nehalem-based systems. This support enables unused CPU resources to be dynamically placed in a state where they consume a fraction of the power consumed in their normal operating state. This feature also provides Solaris support for the power saving feature, as well as the policy implementation that decides when idle CPUs should request deep idle mode. This feature will be enabled by default where supported, and can be disabled through the `cpu-deep-idle` keyword in `power.conf(4)`.

- Observability for Intel's Turbo Mode feature – Intel Nehalem-based systems have the ability to raise the operating frequency of a subset of the available cores when there is enough thermal headroom to do so. This ability temporarily boosts performance, but it is controlled by the hardware and transparent to software. Starting with the Solaris 10 5/09 release, a new `kstat` module observes when the system is entering the turbo mode and at which frequency it operates.

Developer Tools Enhancements

The following developer tools features and enhancements have been added to the Solaris 10 5/09 release.

SunVTS 7.0 Patch Set 5

SunVTS™ 7.0 Patch Set 5 has the following enhancements:

- Infrastructure enhancements:
 - Ability to specify device specific options in a test
 - Creation of generic or host-specific sessions for testing
 - Loop function on a particular test pass
 - Terminal user interface (UI) support for the creation or loading of generic and host-specific sessions
- CPU Diagnostics Enhancements:
 - `System test`, `sys test`, performs processor-level isolation in the event of a failure
 - `CPU test`, `cpustest`, is a multiprocess test. A single test binary can test all the CPUs in the system simultaneously.
- Memory Diagnostics Enhancements:
 - `physmem`-based `ramtest` has the option to read address length in Kbytes, Mbytes, and Gbytes
 - Improved `l3` buffer test with added memory, march-test algorithms
- IO Diagnostics Enhancements:
 - New `hlgraphics test` test is added for testing graphics cards
 - Users can specify back-to-back loopback option for the `nxge` interface in the network test
 - `Cddvdtest` is enhanced to support different drive speeds
 - `Disktest` is enhanced to support the following features:
 - Tuned to stress USB storage devices
 - Perform disk performance testing
 - Not perform write testing on the root disk
 - Test Solid State Devices (SSD) with test level, wear-leveling mechanism

- Supports read-and-write buffer cache test

x86: CPU Performance Counter Updates for Intel Processors

Modern microprocessors contain hardware performance counters that enable the measurement of many different hardware events related to CPU behavior. Hardware events include instruction and data cache misses as well as various internal states of the processor. Data from the performance counters can be used to analyze and tune the behavior of software on a particular type of processor. The Solaris 10 5/09 OS provides access to CPU Performance Counters (cpc) through the `libcpc(3LIB)` interface and through the `cpurack(1)` and `cpustat(1M)` utilities.

Driver Enhancements

The following driver features and enhancements have been added to the Solaris 10 5/09 release.

hermon Driver

This feature introduces a Solaris driver for the fourth generation of InfiniBand (IB) HCA chips from Mellanox, Ltd. The hermon driver provides IB support for SDR, DDR, and QDR chips for conventional HCAs, EMs, and NEMs for blade environments.

The hermon driver enables higher bandwidth and lower latency in IB transmissions, compared to previous generations of the IB product. The higher bandwidth and lower latency are most important in high-performance computing (HPC) applications, though the increase in performance is advantageous in all environments.

In addition, the uDAPL library, a critical underpinning of the MPI library, is updated to work with this driver, providing optimal performance with MPI-based applications.

iSCSI Target

Starting with the Solaris 10 5/09 release, iSCSI Target is upgraded to provide new features and functionality.

This iSCSI Target update includes the following performance, scalability, interoperability, and reliability improvements:

- Improved TCP/IP timeout recovery
- iSCSI initiator invoked SCSI RESETs
- Code path and memory leak cleanup
- Improved interoperability with Target Port Group Tags (TPGT), unidirectional and bidirectional CHAP authentication, and RADIUS server support
- Improved Internet Storage Name Service (iSNS) support, including recovery from unavailable iSNS servers
- Updated SCSI-3 Persistent Reserve functionality that enables the use of the functionality in various clustering solutions on both Solaris and other operating systems

The Solaris iSCSI Target release now supports a wide variety of iSCSI initiators for the following operating systems:

- Solaris 10
- OpenSolaris
- Linux: Red Hat Enterprise Linux (RHEL), Suse, and Ubuntu
- VMWare ESX
- Microsoft Windows (XP, Vista, Server 2003, Server 2008, Windows Cluster Server)
- Mac OS X

x86: NetXen 10-GigE Device Driver

The `ntxn(7D)` is a new NIC driver that supports NetXen's PCI Express-based 10-Gigabit Ethernet network interface cards (NIC). Users can access the network through Solaris OS on platforms that have a NetXen NIC installed.

Intel ICH10 and Hartwell NIC Support in E1000g Driver

Starting with the Solaris 10 5/09 release, the ICH10 and Hartwell network interfaces are the default network interface cards (NIC) on some x64 and x86 machines. Users can access the network easily with these network interfaces.

xge Driver Can Enable Multiple Receive Rings and MSI-X

The `xge` driver enables multiple receive rings and MSI-X if the driver can allocate enough MSI-X vectors on platforms that support MSI-X. The performance of the driver is enhanced by this feature. If the driver is unable to allocate enough MSI-X vectors, the driver continues to work as before in the legacy interrupt mode.

Language Support Enhancements

The following language support enhancement has been added to the Solaris 10 5/09 release.

New Locale Support for Kazakhstan and Ukraine

The Solaris 10 5/09 release now supports the Kazakhstan `kk_KZ.UTF-8` and Ukraine `uk_UA.UTF-8` locales.

Additional Software Enhancements

The following additional software feature has been added to the Solaris 10 5/09 release.

SPARC: Fp-scrubber Daemon

The `Fp-scrubber` is a user-level daemon that periodically runs nonintrusive tests to validate proper functioning of the floating-point unit (FPU) hardware. When an error is detected by the test, a fault management action is initiated by using the `fmd(1M)` command. The `Fp-scrubber` daemon supports only UltraSPARC III and UltraSPARC IV class of processors.

What's New in the Solaris 10 10/08 Release

This chapter summarizes new features in the Solaris 10 10/08 release.

Installation Enhancements

The following installation features and enhancements have been added to the Solaris 10 10/08 release.

Solaris Installation for ZFS Root Pools

Starting with the Solaris 10 10/08 release, you can install and boot a ZFS root pool.

The following installation programs perform an initial installation of a ZFS root pool:

- The Solaris text installer performs an initial installation for a ZFS root pool. During the installation, you can choose to install either a UFS file system or a ZFS root pool. You can set up a mirrored ZFS root pool by selecting two disks during the installation. Or, you can attach or add additional disks after the installation to create a mirrored ZFS root pool. Swap and dump devices on ZFS volumes are automatically created in the ZFS root pool.

For step-by-step instructions, see Chapter 3, “Installing With the Solaris Interactive Text Installer for ZFS Root Pools (Planning and Tasks),” in *Solaris 10 Installation Guide: Basic Installations*.

- With custom JumpStart, you can create a profile to create a ZFS storage pool and designate a bootable ZFS file system. New ZFS profile keywords install a ZFS root pool for an initial installation. A ZFS profile contains a limited set of keywords.

For more information about JumpStart and ZFS, see Chapter 9, “Installing a ZFS Root Pool With JumpStart,” in *Solaris 10 Installation Guide: Custom JumpStart and Advanced Installations*.

You can use Solaris Live Upgrade to perform the following tasks:

- Migrate a UFS root (/) file system to a ZFS root pool
- Create a new boot environment in the following ways:
 - Within an existing ZFS root pool
 - Within another ZFS root pool
 - From a source other than the currently running system
 - On a system with non-global zones installed

After you have used the `lucreate` command to create a ZFS boot environment, you can use other Solaris Live Upgrade commands on the boot environment, such as the `luupgrade` and `luactivate` commands. For more information on using Solaris Live Upgrade for ZFS, see Chapter 12, “Solaris Live Upgrade and ZFS (Overview),” in *Solaris 10 10/08 Installation Guide: Solaris Live Upgrade and Upgrade Planning*.

System Administration Enhancements

The following system administration features and enhancements have been added to the Solaris 10 10/08 release.

Solaris Installation Tool Support of ZFS File Systems

The following Solaris installation tools have been enhanced in the Solaris 10 10/08 release to support ZFS file systems:

- Solaris interactive text installer to install a UFS or a ZFS root file system. The default file system is still UFS for the Solaris 10 10/08 release.
- Custom JumpStart features to set up a profile to create a ZFS storage pool and designate a bootable ZFS file system.
- Migrate a UFS root file system to a ZFS root file system by using the Solaris™ Live Upgrade feature. The `lucreate` and `luactivate` commands have been enhanced to support ZFS pools and file systems. The `lustatus` and `ludelete` commands work as in previous Solaris releases.
- Set up a mirrored ZFS root pool by selecting two disks during the installation. Or, you can attach or add additional disks after the installation to create a mirrored ZFS root pool.
- Automatically create swap and dump devices on ZFS volumes in the ZFS root pool.

For more information, see the *Solaris ZFS Administration Guide*.

For previous Solaris 10 releases, see the following limited Solaris installation tool support for ZFS file systems:

- Custom JumpStart – You cannot include ZFS file systems in a JumpStart profile. However, you can run following scripts from a ZFS storage pool to set up an install server or install client:

- `setup_install_server`
- `add_install_server`
- `add_install_client`
- Live Upgrade – Preserves your original boot environment and carries over your ZFS storage pools into the new environment. Currently, ZFS cannot be used as a bootable root file system so your existing ZFS file systems are not copied into the boot environment.
- Solaris Initial Install - ZFS file systems are not recognized during an initial installation. However, if you do not specify any of the disk devices that contain ZFS storage pools to be used for the installation, you should be able to recover your storage pools by using the `zpool import` command after the installation. For more information, see the [`zpool\(1M\)`](#) man page.
As with most reinstallation scenarios, you should back up your ZFS files before proceeding with the initial installation option.
- Solaris Upgrade – Your ZFS file systems and storage pools are preserved.

SunVTS 7.0 Patch Set 3

SunVTS™ 7.0 Patch Set 3 follows a conventional three-tier architecture model that includes a browser-based user interface, a Java™ technology-based middle server, and a diagnostic agent. SunVTS Patch Set 3 has the following enhancements:

- Disk and Network tests provide device selection and deselection capability.
- The browser-based UI and TTY UI provide support for checkboxes in the logical test (LT) options.
- `fptest` and cache tests have been enhanced for x86 platforms.
- `disktest` has been updated and does not perform any write-test on file system partitions.

For more information about SunVTS, see <http://docs.sun.com/app/docs/prod/test.validate>.

lockstat Provider for DTrace

Starting with the Solaris 10 10/08 release, DTrace `lockstat` probes that displayed the spin count (spins) now returns spin time in nanoseconds. The `lockstat` provider interface and the `lockstat` command options have not changed.

System Resource Enhancements

The following system resource features and enhancements have been added to the Solaris 10 10/08 release.

New Solaris Zones Features

The following Solaris Zones enhancements are found in the Solaris 10 10/08 release:

- **Update on Attach** – If the new host has the same or later versions of the zone-dependent packages and their associated patches, using `zoneadm` at `attach` with the `-u` option, updates those packages within the zone to match the new host. If the new host has a mixture of higher and lower version packages and patches as compared to the source host, then an update during the attach operation is not allowed. This option also enables automatic migration between machine classes, such as from `sun4u` to `sun4v`.

For more information, see the [zoneadm\(1M\)](#) man page and *System Administration Guide: Solaris Containers-Resource Management and Solaris Zones*.

- **Ability to Set Default Router in Shared-IP Zone** – A `defrouter` property has been added to the `net` resource in the `zonecfg` utility for shared-IP non-global zones. You can set the default router for the network interface through this property.

For more information, see the [zonecfg\(1M\)](#) man page and *System Administration Guide: Solaris Containers-Resource Management and Solaris Zones*.

- **ZFS Zone Path Permitted** – Starting with the Solaris 10 10/08 release, the `zonepath` can be on ZFS and the system can be upgraded. For a zone with the `zonepath` on ZFS, only Solaris Live Upgrade can be used to upgrade the system. For more information, see *System Administration Guide: Solaris Containers-Resource Management and Solaris Zones*.

x86: New GRUB `findroot` Command

All Solaris installation methods, including Solaris Live Upgrade, now use the `findroot` command for specifying which disk slice to boot on an x86 based system. Previously, the `root` command, `root (hd0.0.a)`, was used to explicitly specify which disk slice to boot. This information is located in the `menu.lst` file that is used by GRUB. The most common form of the entry in the `menu.lst` file is:

```
findroot (rootfs0,0,a)
kernel$ /platform/i86pc/kernel/$ISADIR/unix
module$ /platform/i86pc/$ISADIR/boot_archive
```

The `findroot` command has the ability to discover the targeted disk, irrespective of the boot device. In addition, the `findroot` command provides enhanced support for booting systems with ZFS roots, as well as those with UFS roots. In addition to the `findroot` command, the name of a signature file on the slice, `(<mysign>, 0, a)` is now provided. The boot signature is located in the `/boot/grub/bootsign` directory on the system. The name of the signature file varies, depending on the installation method that is used.

For more information, see “[Implementation of the findroot Command](#)” in *System Administration Guide: Basic Administration*.

x64: Support for 256 Processors

Starting with the Solaris 10 10/08 release, the Solaris 64-bit operating system supports up to 256 processors on the x86 platform. Previous releases of the Solaris OS supported only up to 64 processors on x86 platforms in the 64-bit mode.

Support for 256 processors provides users the following benefits:

- Users can use this release of Solaris on their existing x86 systems with up to 256 processors.
- Users need not upgrade their OS when they upgrade to machines with over 64 processors in the future.

Users using this release of Solaris will be able to retain their entire software stack including the operating system, when they purchase larger machines.

System Performance Enhancements

The following system performance features and enhancements have been added to the Solaris 10 10/08 release.

SPARC: Solaris SPARC Boot Architecture Redesigned

The Solaris SPARC bootstrap process has been redesigned to increase commonality with the Solaris x86 boot architecture.

The improved Solaris boot architecture brings direct boot, ramdisk-based booting, and the ramdisk miniroot to the SPARC platform. These enabling technologies support the following functions:

- Booting a system from additional file system types. For example, a ZFS file system.
- Booting a single miniroot for software installation from DVD, NFS, or HTTP

Additional improvements include significantly faster boot times, increased flexibility, and reduced maintenance requirements.

As part of this architecture redesign, the Solaris boot archives and the `bootadm` command, previously only available on the Solaris x86 platform, are now an integral part of the Solaris SPARC boot architecture.

The primary difference between the SPARC and x86 boot architectures is how the boot device and file are selected at boot time. SPARC-based systems continue to use the OpenBoot PROM (OBP) as the primary administrative interface, with boot options selected by using OBP commands. On the x86 based platform, these options are selected through the BIOS and the GRand Unified Bootloader (GRUB) menu.

Note – In the Solaris 10 10/08 release, the ability to directly load and boot the UNIX kernel is only available on the SPARC platform. The x86 platform continues to use the multiboot style of booting.

Although the implementation of the Solaris SPARC boot has changed, no administrative procedures for booting a SPARC-based system have been impacted. Boot tasks that are performed by the system administrator remain as they were prior the boot architecture redesign.

For more information, see the following:

- [bootadm\(1M\)](#) and [boot\(1M\)](#) man pages
- Chapter 12, “Booting a Solaris System (Tasks),” in *System Administration Guide: Basic Administration* for instructions on booting a Solaris system

x86: Kernel Support for Intel SSSE3, SSE4.1, SSE4.2, and AMD SSE4A

The kernel now detects the presence of existing Intel SSSE3, SSE4.1, SSE4.2, and AMD SSE4A instruction sets. This feature enables loading and executing programs that require these hardware capabilities. [dis\(1\)](#) now supports disassembly of instructions from these instruction set extensions.

For a list of hardware capability requirements for a given binary, see the [file\(1\)](#) command. For a list of hardware capabilities supported by the loader on a given machine, see [isainfo\(1\)](#).

Security Enhancements

The following networking features and enhancements have been added to the Solaris 10 10/08 release.

Separation of Duty Enforcement Through the Solaris Management Console

This feature enforces that two or more people are required to manage users through the Solaris Management Console (SMC). Separation of duty is enforced by rule. The System Administrator role creates users, but cannot assign passwords and rights. The Security Administrator role assigns passwords and rights, but cannot create users.

Separation of Duty is an accreditation requirement for government customers. SMC now supports this feature and makes it easier to achieve security-level certification.

For more information, see “[Create Rights Profiles That Enforce Separation of Duty](#)” in *Solaris Trusted Extensions Configuration Guide*.

SHA256/SHA512 crypt(3C) Plug-in

This feature provides an additional pair of [crypt\(3C\)](#) plug-ins based on the SHA256 and SHA512 digest algorithms. This plug-in provides a [crypt\(3C\)](#) hash that uses FIPS 140-2 approved algorithms and discontinues using MD5–based hashes.

pam_list Module

The [pam_list](#) module provides functions to validate the user's account on a specific host based on a list of users and netgroups. This module can be used as a quick replacement for account validation through the [passwd_compat](#) mode.

For more information, see the [pam_list\(5\)](#) man page.

Desktop Enhancements

The following desktop features and enhancements have been added to the Solaris 10 10/08 release.

SPARC: Adobe Reader 8.1.2

Starting with the Solaris 10 10/08 release, Adobe® Reader 8.1.2 is included with the Solaris OS. The new version of Adobe Reader includes many new features, including a much improved user interface, support for Shared Reviews, and several security fixes.

For more details, see http://blogs.adobe.com/acroread/2008/02/adobe_reader_812_for_linux_and.html.

Flash Player 9.0.124.0

Starting with the Solaris 10 10/08 release, the Solaris OS includes Adobe Flash Player 9.0.124.0. New features in this version of the Flash Player include the following:

- H.264 video and HE-AAC audio codec support
- Full screen mode for the Solaris OS
- Improved performance due to multi-core support, hardware and image scaling, multi-threaded video decoding, Flash Player cache, and Flash Media Server buffering

For more information, see <http://www.adobe.com/products/flashplayer/productinfo/features/>.

Networking Enhancements

The following networking features and enhancements have been added to the Solaris 10 10/08 release.

Communication Protocol Parser Utilities

Session Description Protocol (SDP), RFC 4566, is used for describing multimedia sessions for session announcement, session invitation, and other forms of multimedia session initiation. SDP conveys media details such as type and encoding, transport protocol, session name, purpose, owner, and other session description metadata to the participants.

The `libcommputil(3LIB)` library provides public interfaces that parses the SDP description and checks for syntax conformance. This library also contains interfaces to generate SDP messages and convert the messages to byte-strings. SDP is used predominantly by the Session Initiation Protocol (SIP). With the `libcommputil(3LIB)` library, Solaris SIP developers can leverage these interfaces in developing SIP applications on the Solaris platform.

For more information, see the `libcommputil(3LIB)` man page.

SIP End-to-end Traffic Measurements and Logging

The Solaris SIP stack now provides SIP application developers with the following two new features:

- **End-to-end traffic measurements** – The end-to-end traffic measurement feature tracks the following activities.
 - Total number of bytes sent and received by the stack
 - Total number of SIP requests and responses sent and received
 - Number of SIP requests sent and received broken down by methods
 - Number of SIP responses sent and received broken down by response codes
- **SIP dialog or transaction logging** – The SIP logging feature can enable and disable dialog or transaction logging. All SIP messages that are exchanged within a dialog or transaction are captured and stored in a log file provided by the application, when the dialog or transaction is terminated. The SIP logging feature helps developers in call tracing and debugging.

Device Management Enhancements

The following device management features and enhancements have been added to the Solaris 10 10/08 release.

Faulty Device Retirement Feature

Starting with the Solaris 10 10/08 release, the Solaris OS includes a new device retirement mechanism to isolate a device as *faulty* by the fault management framework (FMA). This feature allows faulty devices to be safely and automatically inactivated to avoid data loss, data corruption, panics, and system down time. The retirement process is done safely, taking into account the stability of the system after the device has been retired.

Critical devices are never retired. If you need to manually replace a retired device, use the `fmadm repair` command after the device replacement so that system knows that the device is replaced, in addition to the manual replacement steps.

The `fmadm repair` process is as follows:

- Identify the faulty device with the `fmadm faulty -a` command.

```
# fmadm faulty
```

```
STATE RESOURCE / UUID
```

```
-----
```

```
faulty <fmri>
```

- Clear the fault by using the `fmadm repair` command.

```
# fmadm repair <fmri>
```

- Run the `fmadm faulty` command again to be sure the fault is cleared.

```
# fmadm faulty -a
STATE RESOURCE / UUID
```

For more information, see [fmadm\(1M\)](#).

A general message regarding device retirement is displayed on the console and written to the `/var/adm/messages` file so that you are aware of a retired device. For example:

```
Aug 9 18:14 starbug genunix: [ID 751201 kern.notice]
NOTICE: One or more I/O devices have been retired
```

You can use the `prtconf` command to identify specific retired devices. For example:

```
# prtconf
.
.
.
pci, instance #2
    scsi, instance #0
        disk (driver not attached)
        tape (driver not attached)
    sd, instance #3
    sd, instance #0 (retired)
    scsi, instance #1 (retired)
        disk (retired)
        tape (retired)
pci, instance #3
    network, instance #2 (driver not attached)
    network, instance #3 (driver not attached)
os-io (driver not attached)
iscsi, instance #0
pseudo, instance #0
.
.
.
```

MPxIO Support for Hitachi Adaptable Modular Storage Series Arrays

Starting with the Solaris 10 10/08 release, the Hitachi Adaptable Modular Storage (AMS) and Hitachi Workgroup Modular Storage systems are integrated with MPxIO. Dual controller Adaptable Modular Storage and Workgroup Modular Storage models can fully utilize multiple paths in MPxIO environments. This support also enables the full functionality of MPxIO with the Hitachi storage systems in Sun Cluster environments.

Driver Enhancements

The following driver features and enhancements have been added to the Solaris 10 10/08 release.

x86: NVIDIA ck804/mcp55 SATA Controller Driver

`nv_sata` is a SATA HBA driver capable of hot-pluggable functions, for NVIDIA ck804/mcp55 and compatible SATA controllers.

For more information, see the [nv_sata\(7D\)](#) man page.

x86: LSI MegaRAID SAS Controllers Driver

The new LSI `mega_sas` driver supports the following controllers:

- Dell PERC 5/E, 5/i, 6/E, and 6/i RAID Controllers
- IBM ServeRAID-MR10k SAS/SATA Controller
- LSI MegaRAID SAS 8308ELP, 8344ELP, 84016E, 8408ELP, 8480ELP, 8704ELP, 8704EM2, 8708ELP, 8708EM2, 8880EM2, and 8888ELP Controllers

The `mega_sas` driver supports the following RAID features:

- RAID levels 0, 1, 5, and 6, and RAID spans 10, 50, and 60
- Online capacity expansion (OCE)
- Online RAID Level Migration (RLM)
- Auto resume after loss of system power during array rebuild or reconstruction (OCE/RLM)
- Configurable stripe size up to 1 Mbyte
- Consistency check for background data integrity
- Patrol read for media scanning and repairing
- 64 logical drive support
- Up to 64TB logical unit number (LUN) support
- Automatic rebuild
- Global and dedicated hot spare support

For more information about the LSI MegaRAID products, see http://www.lsi.com/storage_home/products_home/internal_raid/megaraid_sas/index.html website.

ixgbe Driver

Starting with the Solaris 10 10/08 release, the `ixgbe` driver is integrated with the Solaris OS. The `ixgbe` is a 10 Gigabit PCI Express Ethernet driver that supports Intel 82598 10 Gigabit Ethernet controller.

SPARC: Support for aac Driver

Starting with the Solaris 10 10/08 release, the HBA driver for Adaptec Advanced RAID Controller, aac, now supports the SPARC platform.

For more information, see the [aac\(7D\)](#) man page.

Additional Software Enhancements

The following additional software features and enhancements have been added to the Solaris 10 10/08 release.

Perl Database Interface and Perl PostgreSQL Driver

Perl Database Interface (DBI) is a generic database interface to talk to specific database back-end. DBD::Pg is a PostgreSQL driver which enables Perl applications to interact with PostgreSQL database through DBI.

For more information, see the following:

- <http://dbi.perl.org>
- <http://search.cpan.org/~timb/DBI-1.58/DBI.pm>

PostgreSQL 8.3

PostgreSQL is an advanced, open-source Relational Database Management System (RDBMS). The important features of PostgreSQL 8.3 version are integrated text search, XML support, and performance improvements in many areas.

For more information, see the community web site, <http://www.postgresql.org/>.

Language Support Enhancements

The following language support feature has been added to the Solaris 10 10/08 release.

IIIMF Hangul Language Engine

The Hangul LE (Language Engine) is a new Korean input method that enhances user experience. Hangul LE has the following features:

- User-friendly GUI
- More convenient Hangul or Hanja input functionalities

For more information, see the Hangul LE help.

Freeware Enhancements

The following freeware features and enhancements have been added to the Solaris 10 10/08 release.

C-URL - The C-URL Wrappers Library

C-URL is a utility library that provides programmatic access to the most common Internet protocols such as, HTTP, FTP, TFTP, SFTP, and TELNET. C-URL is also extensively used in various applications.

For more information, see <http://curl.haxx.se/>.

Libidn - Internationalized Domain Library

Libidn provides implementations of the Stringprep (RFC 3454), Nameprep (RFC 3491), Punycode (RFC 3492), and IDNA (RFC 3490) specifications. This library provides new functionality and facilities to the Solaris OS.

For more information, see the following resources:

- <http://www.gnu.org/software/libidn/>
- <http://www.ietf.org/rfc/rfc3490.txt>
- <http://www.ietf.org/rfc/rfc3492.txt>
- <http://www.ietf.org/rfc/rfc3490.txt>
- <http://www.unicode.org/reports/tr28/tr28-3.html>

LibGD - The Graphics Draw Library

LibGD is a graphics conversion and manipulation utility library. This library is used extensively in web-based application frameworks. The command-line utilities of LibGD provide easy-to-use graphics-conversion facilities.

For more information, see <http://www.libgd.org/>.

TIDY HTML Library

TIDY is an HTML parser. This parser is the HTML equivalent of `lint(1)`. TIDY is useful in validating the accuracy of static and dynamic HTML pages.

For more information, see <http://tidy.sourceforge.net/>.

What's New in the Solaris 10 5/08 Release

This chapter summarizes all the new features in the Solaris 10 5/08 release.

System Administration Enhancements

The following system administration features and enhancements have been added to the Solaris 10 5/08 release.

Solaris Trusted Extensions Administrator Procedures

Starting with this release, the Solaris™ Trusted Extensions packages are installed when the Solaris OS is installed. The `ExtraValue` directory is no longer present. This directory previously included the Solaris Trusted Extensions packages. The Solaris Trusted Extensions functionality is managed by the service management facility (SMF) as the `svc:/system/label:default` service. This service must be enabled. After the service is in the online state, reboot the system to activate Solaris Trusted Extensions. Additional configuration is required after the reboot. For more information, see *Solaris Trusted Extensions Configuration Guide*.

The Solaris 10 5/08 release also includes the following features:

- The domain of interpretation (DOI) is configurable. For more information, see “Network Security Attributes in Trusted Extensions” in *Solaris Trusted Extensions Administrator’s Procedures*.
- The NFSv3 protocol supports multilevel mounts. For more information, see “Trusted Extensions Software and NFS Protocol Versions” in *Solaris Trusted Extensions Administrator’s Procedures*.
- The name service cache daemon, `nscd`, can be separately configured per labeled zone. This configuration supports environments where each zone is connected to a subnetwork that runs at the label of the zone, and the subnetwork has its own name server for that label.

For more information about Solaris Trusted Extensions, see *Solaris Trusted Extensions Administrator’s Procedures*.

Flash Update Tool

`fwflash(1M)` is a new Solaris command for the manipulation of firmware for PCI-X, and PCI-Express HBA and HCA cards. Currently, the command enables listing, reading, and writing the firmware for the InfiniBand HCA cards.

For more information about this command, see the [fwflash\(1M\)](#) man page.

PPD File Management Utility

The PostScript™ Printer Description (PPD) file management utility, `/usr/sbin/ppdmgr`, manages PPD files that are used with the Solaris print subsystem.

By using the `ppdmgr` utility, you can perform the following tasks:

- Add a PPD file to a PPD file repository on a system
- Supply a label to group PPD files within a PPD file repository
- Update the cache of the PPD file information that is used by the Solaris Print Manager (`printmgr`) GUI to display supported printer information

You can add a new PPD file by using the `ppdmgr` utility or by using the `lpadmin -n` command. When you add a new PPD file, you automatically update the cache of the PPD file information that the `printmgr` GUI uses to display supported printer information.

Note – The delivery location of PPD files in the Solaris OS has changed. During a software upgrade, any print servers with print queues that were defined by using PPD files from the previous PPD file delivery location are automatically updated to reflect the new PPD file delivery location.

In addition, a new SMF service, `print/ppd-cache-update`, has been introduced. The `print/ppd-cache-update` service is enabled by default. This service runs one time during system reboot to update the printer cache information with changes from all the PPD file repositories on the system. The service might take longer to transition to an online state during a system reboot after a software installation or upgrade. Also, if any changes were made to the PPD file repositories since the last PPD cache update, during system reboot, the service might take longer to come online. Changes made to the PPD file repositories on a system are not reflected in the PPD cache used by Solaris Print Manager until the `print/ppd-cache-update` service is online.

For more information, see the following:

- [ppdmgr\(1M\)](#) man page
- Chapter 9, “Administering Printers by Using the PPD File Management Utility (Tasks),” in *System Administration Guide: Solaris Printing*

Internet Printing Protocol Client-Side Support

Client-side support for the Internet Printing Protocol (IPP) enables Solaris client systems to communicate with IPP-based print services, such as those on the Linux and Mac OS X operating systems, as well as other platforms.

Small improvements are also featured in the server-side support for the IPP listening service. These improvements promote better interoperability, including some minor changes that result in a more standard representation of printer and job attribute data.

The IPP server and client implementation in the Solaris OS is one of several OpenSolaris™ printing projects that are currently under development. OpenSolaris printing provides a set of specifications and implementations of software that enables you to create standardized, scalable printing components for the Solaris and Linux software, or any operating system that contains a set of POSIX interfaces.

For more information, see the *System Administration Guide: Solaris Printing*.

For more information about OpenSolaris Printing, see <http://opensolaris.org/os/community/printing/>.

Selectable Use of localhost for Solaris Print Server Database Hostname

This printing feature enables the Solaris print system to recognize and use localhost as the local host in the print system databases. In prior releases, /bin/hostname was used solely to generate the print hostname. The print system depended on this name remaining constant. The ability to use localhost as the name of the current system enables print servers to maintain the same print hostname, independent of the system's host name.

Note – The modification applies to the setup of local print queues exclusively.

To support this feature, the following modifications are effective for the `lpadmin` command and the Solaris Print Manager graphical user interface (GUI):

- The `lpadmin` command uses the `-s` option when creating a local print queue.

To use localhost as the host name that is specified within the print server, set the print hostname to localhost, as shown:

```
# lpadmin -p <new-print-queue> -s localhost -v <device>
```

For example:

```
# lpadmin -p foo -s localhost -v /dev/term/a
```

Note – The default behavior of the `lpadmin` command has not changed.

- Solaris Print Manager now includes an added tool attribute check box, Use `localhost` for Printer Server. The `localhost` attribute is selected by default. To deselect the `localhost` attribute, uncheck the box. Unchecking the box selects the previously chosen behavior for this attribute.

For more information, see the following:

- [printmgr\(1M\)](#) man page
- [lpadmin\(1M\)](#) man page
- *[System Administration Guide: Solaris Printing](#)*

Fault Management for T5140/T5240 Platforms

The Solaris predictive self-healing technology is available on Sun SPARC Enterprise T5140 and T5240 platforms. Predictive self-healing features include the following:

- Automated error handling
- Automated diagnosis
- Automated recovery for CPU, memory and I/O subsystems
- Clear and concise error messages

For more information, see http://www.sun.com/software/solaris/ds/self_healing.jsp and <http://opensolaris.org/os/community/fm>.

SunVTS 7.0

SunVTS™ is a comprehensive system validation and test suite designed to support Sun hardware platforms and peripherals. SunVTS 7.0 is the next generation of SunVTS 6.0 and its compatible versions.

SunVTS 7.0 includes the following features:

- Introduction of the concept of purpose-based testing
- Improved diagnostics effectiveness
- Web-based user interface
- Simplified usage
- New architecture framework
- Enterprise View

SunVTS 7.0 follows a conventional three-tier architecture model. This model is composed of a browser-based user interface, a Java based middle server, and a diagnostic agent.

System Resource Enhancements

The following system resource features and enhancements have been added to the Solaris 10 5/08 release.

Solaris Zones and Branded Zones

You can perform a trial run of a zone migration before the zone is moved to the new machine by using the “no execute” option, `-n`.

The `zoneadm detach` subcommand is used with the `-n` option to generate a manifest on a running zone without actually detaching the zone. The state of the zone on the originating system is not changed. The `zoneadm attach` subcommand is used with the `-n` option to read this manifest and verify that the target machine has the correct configuration to host the zone without actually doing an attach.

For more information, see the following:

- Chapter 23, “Moving and Migrating Non-Global Zones (Tasks),” in *System Administration Guide: Solaris Containers-Resource Management and Solaris Zones*
- [zoneadm\(1M\)](#)

CPU Caps

CPU caps provide absolute fine-grained limits on the amount of CPU resources that can be consumed by a project or a zone. CPU caps are provided as a `zonecfg` resource, and as project and zone-wide resource controls.

- The `zonecfg capped-cpu` resource provides an absolute limit on the amount of CPU that can be consumed by a project or a zone.
- The following resource controls are available:

<code>zone.cpu-cap</code>	Absolute limit on the amount of CPU resources that can be consumed by a non-global zone.
<code>project.cpu-cap</code>	Absolute limit on the amount of CPU resources that can be consumed by a project.

For more information, see the following:

- [zonecfg\(1M\)](#) man page
- [zones\(5\)](#) man page
- *System Administration Guide: Solaris Containers-Resource Management and Solaris Zones*

projmod(1M) Option

Use the `projmod` command with the `-A` option to apply the resource control values defined in the project database to the active project. Existing values that do not match the values defined in the project file, such as values set manually by `prctl(1)`, are removed.

Device Management Enhancements

The following device management features and enhancements have been added to the Solaris 10 5/08 release.

Tape Self-Identification

The tape self-identification feature configures the tape automatically, with the parameters provided by the tape drive. Previously, the configuration data for a tape drive was statically supplied through user-editable configuration files, built-in configuration tables, or default values. The tape-self identification feature uses a few SCSI commands to directly query the required parameters from the tape drive. When the `st` driver gets the parameters, the tape drive uses them on the Solaris OS.

The advantages of tape-self identification over the traditional file-based configuration are:

- Simple and no user intervention is needed
- Faster support for new tape drives
- Easy to use, standard-based interface

x86: Enhanced Speedstep CPU Power Management

Starting with this release, Intel's Enhanced Speedstep™ technology is supported on the Solaris OS. Enhanced Speedstep support enables Solaris platform users to manage the power consumption of their Intel processors by lowering the processor frequency during idle periods.

For more information about how to enable Solaris CPU power management, see the [power.conf\(4\)](#) man page.

x86: PowerNow! CPU Performance Management

Starting with this release, AMD's PowerNow! technology is supported on the Solaris OS. PowerNow! support enables Solaris platform users to manage the power consumption of their Opteron 10h family of processors by adjusting the processor operating frequency and voltage according to the task being performed.

For more information about how to enable Solaris CPU power management, see the [power.conf\(4\)](#) man page.

iSNS Support in the Solaris iSCSI Target

This Solaris release provides support for the Internet Storage Name Service (iSNS) protocol in the Solaris iSCSI target software. The iSNS protocol enables automated discovery, management, and configuration of iSCSI devices on a TCP/IP network.

The Solaris iSCSI target software does not include native iSNS server support. However, in this Solaris 10 release, you can add access to an existing iSNS server to automatically discover the iSCSI devices in your network.

The `iscsitadm` command is used to configure the Solaris iSCSI target to discover the iSNS server and enable or disable the iSNS discovery. Use the hostname or the IP address to specify the iSNS server.

For more information, see the [iscsitadm\(1M\)](#) man page and [Chapter 14, “Configuring Solaris iSCSI Targets and Initiators \(Tasks\)”](#) in *System Administration Guide: Devices and File Systems*.

Security Enhancements

The following security features and enhancements have been added to the Solaris 10 5/08 release.

Solaris Trusted Extensions Supports Mounting Labeled File Systems With the NFSv3 Protocol

Starting with this release, the Solaris Trusted Extensions software can mount labeled file systems by using NFS Version 3 (NFSv3) in addition to NFS Version 4 (NFSv4). Solaris Trusted Extensions has no restrictions on using TCP as an underlying transport protocol for NFS. However, users cannot choose User Datagram Protocol (UDP) as the underlying protocol for read-down NFS access for NFSv3. The use of UDP for the initial mount operation is supported, but UDP is not supported for subsequent multilevel NFSv3 operations.

SPARC: Hardware -Accelerated Elliptical Curve Cryptography (ECC) Support

The UltraSPARC T2 (Niagara 2) based platforms support hardware acceleration of Elliptical Curve Cryptography (ECC) algorithms. The Solaris OS now supports high performance ECDSA and ECDH on these platforms. These new ECC algorithms are accessible to all users of the Solaris Cryptographic Framework including Java technology and OpenSSL users.

Networking Enhancements

The following networking features and enhancements have been added to the Solaris 10 5/08 release.

Sockets Direct Protocol

The Sockets Direct Protocol (SDP) is a transport protocol layered over the Infiniband Transport Framework (IBTF). SDP is a standard implementation based on Annex 4 of the Infiniband Architecture Specification Vol1. SDP provides reliable byte-stream, flow controlled two-way data transmission that is very similar to TCP.

For more information see [sdp\(7D\)](#) man page.

inetd Backlog Queue Size

Starting with this release, a tunable to set the backlog queue size of the inetd managed services is introduced. This feature adds an SMF property to inetd called `connection_backlog` that enables the queue size to be modified. The default value of the `connection_backlog` queue size is 10. You can modify the `connection_backlog` property by using the `inetadm` command. For example:

- To list the properties, type:

```
#inetadm -l fmri/pattern
```

- To change the value for a specific service, type:

```
#inetadm -m fmri/pattern connection_backlog=new value
```

- To change the value globally, type:

```
#inetadm -M connection_backlog=newvalue
```

For more information, see the [inetadm\(1M\)](#) man page.

X11 Windowing Enhancements

The following X11 windowing features and enhancements have been added to the Solaris 10 5/08 release.

Xvnc Server and Vncviewer Client

VNC provides a remote desktop session over the Remote Frame Buffer (RFB) protocol. RFB clients, better known as VNC viewers, are available for most platforms, in both open-source and commercial releases.

The Solaris 10 5/08 release now includes Xvnc. Xvnc is an X server that is based on the open-source releases from the RealVNC project and X.Org Foundation. Xvnc is displayed to an RFB protocol client over the network without requiring an existing X server session display on the local video hardware. This release also includes RealVNC's vncviewer RFB client to connect to remote VNC servers, and several associated programs for managing these servers.

For more information, see *System Administration Guide: Virtualization Using the Solaris Operating System*. See also the `Xvnc(1)` and `vncviewer(1)` man pages.

Desktop Tools Enhancements

The following desktop tools features and enhancements have been added to the Solaris 10 5/08 release.

StarOffice 8

Starting with this release, StarOffice has been enhanced to the latest version, StarOffice 8.

For more information about StarOffice, see http://www.sun.com/software/star/staroffice/whats_new.jsp.

Flash Player 9

Starting with this release, the Solaris OS includes the Adobe Flash Player 9. For more information about this Flash Player, see <http://www.adobe.com/products/flashplayer/productinfo/features/>.

Pidgin 2.0

Pidgin is a popular open-source instant messaging client. Pidgin 2.0 includes the following features:

- Many improvements to the UI modules including status system, Buddy List, Conversation, and the chat window
- New Yahoo features including Stealth Settings, Doodle, and the `/list` command
- Improved AIM and ICQ file transfers
- Improved Log Viewer module
- Support for the new version of ICQ file transfer
- New IRC features including SSL support, and the new commands `/whowas`, `/nickserv`, `/memoserv`, `/chanserv`, and `/operserv`
- Jabber features including support for SRV lookups, buddy icons, and Jabber User Directory searching

PAPI Print Commands

The Free Standards Group (FSG) Open Printing API (PAPI) commands replace several commonly used print commands, which include the following:

- `cancel(1)`
- `disable(1)`
- `enable(1)`
- `lp(1)`
- `lpstat(1)`
- `lpc(1B)`
- `lpq(1B)`
- `lpr(1B)`
- `lprm(1B)`
- `accept(1M)`
- `lpmove(1M)`

- `reject(1M)`

The implementations of the Open Printing API commands are layered on top of the Free Standards Group Open Printing API in the Solaris OS. This implementation enables the commands to run on top of multiple protocols or services.

Some advantages of the new print command implementations include the following:

- Improved consistency between desktop applications and command-line interfaces
- Multiple print protocols and service support from the command line
- Internet Print Protocol (IPP) client-side support for improved interoperability with Linux, Mac OS X, and other IPP-based print services
- Enhanced remote capability and data when using IPP between print client and server
- The capability to disable network services and retain access to local printers

For more information about the PAPI print commands, see the following:

- “How the PAPI Is Implemented in the Solaris OS” in *System Administration Guide: Solaris Printing*
- OpenSolaris Printing Community web pages at http://opensolaris.org/os/community/printing/projects/papi_client_commands/

System Performance Enhancements

The following system performance features and enhancements have been added to the Solaris 10 5/08 release.

64-bit SPARC: Memory Placement Optimization Support for sun4v Platforms

Memory Placement Optimization (MPO) enables operating systems to allocate memory local to the core where the threads or processes are being executed. The sun4v architecture runs on virtualized hardware environment. The MPO for sun4v platforms feature provides the required standard accessors in the sun4v layer to provide locality information for the generic MPO framework. This feature is effective on platforms that have multiple sockets with differences in memory access latency. The MPO feature enhances the performance of various applications by enabling the OS to allocate memory local to the nodes.

SPARC: Shared Contexts Support

The context mechanism, which is used by the Memory Management Unit (MMU) hardware to distinguish between the use of the same virtual address in different process address spaces, introduces some inefficiencies when shared memory is used. The inefficiencies in shared memory are because the data at a particular shared memory and the address in different processes might really be identical, but the context number associated with each process is different. Therefore, the MMU hardware cannot recognize a match. This inability to recognize

a match results in mappings being unnecessarily evicted from the MMU translation cache and the Translation Lookaside Buffer (TLB), only to be replaced by identical mappings with a different context number.

The Niagara 2 system has an additional shared context, which is a hardware feature that can be used to prevent the inefficiency in handling shared memory. Searching the TLB for mapping a match on either the private or the shared context results in a TLB hit. The current software support for shared context activates the feature for processes that use the Dynamic Intimate Shared Memory (DISM). In this case, the process text segment and DISM segments mapped at the same virtual address with the same permissions for each process use the shared context.

x86: CPUID-Based Cache Hierarchy Awareness

Modern Intel processors provide an interface for discovering information about the cache hierarchy of the processor through the CPUID instruction.

Language Support Enhancements

The following language support features and enhancements have been added to the Solaris 10 5/08 release.

Locale Creator

Locale Creator is a command line and graphical user interface tool that enables users to create and install Solaris locales. Using locale creator users can create installable Solaris packages containing customized locale data of a specific locale. After the created package has been installed, the user has a fully working locale on the system.

For more information, see the following:

- The command `/usr/bin/localectr -h`
- `localectr` man page
- http://developers.sun.com/global/products_platforms/solaris/reference/techart/locale-creator.html

libchewing 0.3.0

Chewing input method (IM) is based on libchewing, which is an open-source library for Traditional Chinese input. libchewing has been upgraded to the libchewing 0.3.0 version. Some of the features of the new version include the following:

- Incompatibility with API/ABI.
- UTF-8 based language engine core for common Unicode environment.
- Includes the libchewing-data subproject.
- Zuin fixes and symbol improvements.

- New binary form of user hash data to speed up loading and solving hash data corruption.
- Improved calculation of internal tree and phone constants.
- Revised tsi.src for richer phrases and avoiding crashes.
- Merge phone and phrase from CNS11643.
- Improved Han-Yu PinYin to use table-lookup implementation.
- Experimental frequency evaluation that re-computes chewing lifetime.
- Implementation of the choice mechanism for symbol pairs.
- Experimental, memory mapping-based, binary data handling to speed up data loading.

For further information, see the *International Language Environments Guide*.

File Encoding Examiner

The File Encoding Examiner (fsexam) enables you to convert the name of a file, or the contents of a plain-text file, from a legacy character encoding to UTF-8 encoding. New features in the fsexam utility include the following:

- Encoding list customization
- Encoding auto-detection
- Support for dry runs, logs, batch conversion, file filtering, symbolic files, command line, and special file types like compress file

For more information, see the fsexam(1) and fsexam(4) man pages.

Kernel Functions Enhancements

The following kernel functions features and enhancements have been added to the Solaris 10 5/08 release.

x86: MONITOR and MWAIT CPU Idle Loop

The Solaris OS uses the SSE3 MONITOR and MWAIT instructions in the x86 processor idle loop. Using the SSE3 instructions in the processor idle loop eliminates the overhead of sending and receiving an interrupt to wake up a halted processor. MONITOR is used to specify a memory range to monitor the idle loop. MWAIT halts the processor until the address previously specified with MONITOR is accessed. With the new idle loop, a processor has to write to memory only to wake up a halted processor.

Driver Enhancements

The following driver features and enhancements have been added to the Solaris 10 5/08 release.

x86: Support Sun Fire X4540 Disk Status Indicators

Starting with this release, the Sun Fire X4540 disk status indicators are supported. The amber Fault status LED and blue Ready to Remove status LEDs are enabled by this feature.

For more information, see the *Sun Fire X4540 Server Diagnostics Guide*.

MPxIO Extension for Serial Attached SCSI Devices on mpt(7D)

The mpt driver has been enhanced to support MPxIO with supported storage devices. When MPxIO is enabled for Serial Attached SCSI (SAS) and SATA devices, they are enumerated under `scsi_vhci(7D)` just like fibre channel devices under `fp(7D)`.

Starting with this release, `stmsboot(1M)` has also been enhanced to support multipathed SAS devices. `stmsboot(1D)` operates on all attached and multipath-capable controllers by default.

If you wish to only enable multipathing on `fp` or `mpt` controllers then you can use the new flag which has been added to restrict operations. The command, `/usr/sbin/stmsboot -D mpt -e`, will enable MPxIO only on attached `mpt` controllers. Replacing `mpt` with `fp` in this command will make `stmsboot` enable MPxIO only on attached `fp` controllers.

x86: SATA ATAPI Support in AHCI Driver

The AHCI driver supports SATA ATAPI CD or DVD devices. Users can use the SATA CD or DVD in AHCI mode instead of the compatible mode. The AHCI mode has better error handling and hot-pluggable capabilities.

For more information, see the [ahci\(7D\)](#) man page.

x86: AMD-8111

The AMD-8111 HyperTransport I/O hub includes a 10/100 Mbps Ethernet LAN controller. The driver is used by the Andretti platform.

SATA NCQ Support in AHCI Driver

The AHCI driver supports the SATA NCQ feature. NCQ support improves performance of the driver.

For more information, see the [ahci\(7D\)](#) man page.

x86: bnx II Ethernet Driver

Starting with this release, support is provided for the Broadcom NetXtreme (bnx) II Ethernet chipset, which includes BRCM5706C, BRCM5706S, BRCM5708C, and BRCM5708S.

For more information, see the [bnx\(7D\)](#) man page.

USB-to-Serial Driver for Keyspan Adapters

Starting with this release, a new driver is provided for Keyspan USB-to-serial adapters. This driver supports the USA-19HS model. This feature enables you to choose between Edgeport adapters and Keyspan adapters.

For further information, see the [usbksp\(7D\)](#) man page.

Freeware Enhancements

The following freeware features and enhancements have been added to the Solaris 10 5/08 release.

32-bit: pgAdmin III

pgAdmin III is a popular and feature-rich, open-source administration and development platform for PostgreSQL. The graphical interface supports all PostgreSQL features and makes administration easy. This tool enables users to write simple SQL queries and also to develop complex databases.

For more information, see <http://www.pgadmin.org/>.

p7zip

Starting with this release, the Solaris OS includes p7zip port. p7zip is similar to the Windows compression and archiving utility, 7zip.

For more information, see <http://p7zip.sourceforge.net/>.

What's New in the Solaris 10 8/07 Release

This chapter summarizes all the new features in the Solaris 10 8/07 release.

System Administration Enhancements

The following system administration features and enhancements have been added to the Solaris 10 8/07 release.

Name Service Switch Enhancements

Enhancements have been made to the name service switch (nss) and to the Name Switch Cache Daemon (nscd(1M)) in order to deliver new functionality. These enhancements include the following:

- Better caching in nscd(1M) and management of connections within the updated framework
- Name service lookups that are access controlled at the naming service on a per-user basis. The updated switch framework adds support for this style of lookups using SASL/GSS/Kerberos in a manner that is compatible with the authentication model used in the Microsoft Active Directory.
- A framework for the future addition of putXbyY interfaces.

For more information about per-user lookups, see *System Administration Guide: Naming and Directory Services (DNS, NIS, and LDAP)*.

iostat Improvements

The -Y option of the iostat command provides new performance information for machines that use Solaris I/O multipathing.

For more information, see the `iostat(1M)` man page.

Solaris System Registration

Starting with this release, you can register the Solaris OS by using one of the following methods:

- Basic Registration 1.1 - Use this method if you want to use Sun Connection's hosted deployment architecture or Update Manager.
- Solaris Registration - Use this method if you want to use Sun Connection to maintain an inventory of systems that you have registered.

Basic Registration 1.1 is a system administration feature that was introduced in the Solaris 10 6/06 release. The Basic Registration feature enables you to create a registration profile and ID to automate your Solaris 10 software registrations for the Update Manager. The Update Manager is the single system update client that is used by Sun Connection. Sun Connection was formerly known as Sun Update Connection System Edition. The Basic Registration wizard appears on system reboot. For information on the Basic Registration 1.1 feature, see “[Basic Registration 1.1](#)” on page 129. For information about Sun Connection's product portfolio and how to register with the wizard, see the Sun Connection Information Hub at <http://www.sun.com/bigadmin/hubs/connection/>.

Solaris Registration enables you to register one or more instances of your Solaris software at the same time by providing a Sun Online Account user name and password. To register, go to <https://sunconnection.sun.com>.

Sun Service Tag

A Sun Service Tag is a product identifier that is designed to automatically discover your Sun systems, software, and services for quick and easy registration. A service tag uniquely identifies each tagged asset, and enables the asset information to be shared over a local network in a standard XML format.

Service tags are enabled as part of the Service Management Facility (SMF) and the SMF `generic_open.xml` profile. If you select the SMF `generic_limited_net.xml` profile, service tags are not enabled.

For more information about SMF, see the *System Administration Guide: Basic Administration*. For more information about service tags, the types of information collected, and automatic registration, see Sun Connection on BigAdmin at <http://www.sun.com/bigadmin/hubs/connection/tasks/register.jsp>.

MPxIO Path Steering

The MPxIO path steering feature includes a mechanism for issuing SCSI commands to an MPxIO LU to be delivered down a specified path to the LU. In order to provide this functionality, a new IOCTL command, `MP_SEND_SCSI_CMD`, is added and is referenced through the existing `scsi_vhci` IOCTL interface. An extension is introduced to the multipath management library (MP-API) which provides access to this new IOCTL command. This enables network administrators to run diagnostic commands through a specified path.

raidctl

`raidctl` is a utility that can perform RAID configuration work by using multiple RAID controllers. The `raidctl` feature contains more detailed information about RAID components, including controller, volume and physical disks. The `raidctl` utility enables the user to track the RAID system more closely and simplify the learning effort on diverse RAID controllers.

For more information, see the following:

- [raidctl\(1M\) man page](#)
- http://www.lsi.com/storage_home/products_home/host_bus_adapters/index.html

Brand-Specific Handlers for `zoneadm` Commands

The `zoneadm(1M)` command is modified to call an external program that performs validation checks against a specific `zoneadm` operation on a branded zone. The checks are performed before the specified `zoneadm` subcommand is executed. However, the external brand-specific handler program for `zoneadm(1M)` should be specified by the brand's configuration file, `/usr/lib/brand/<brand_name>/config.xml`. The external program is specified by the brand's configuration file by using the `<verify_admin>` tag.

To introduce a new type of branded zone, and list brand-specific handlers for the `zoneadm(1M)` subcommand, add the following line to the brand's `config.xml` file:

```
<verify_admin><absolute path to external program> %z %* %*</verify_admin>
```

In this line, `%z` is the zone name, the first `%*` is the `zoneadm` subcommand, and the second `%*` is the subcommand's arguments.

This feature is useful when a given branded zone might not support all the `zoneadm(1M)` operations possible. Brand-specific handlers provide a way to gracefully fail unsupported `zoneadm` commands.

Ensure that the handler program that you specify recognizes all `zoneadm(1M)` subcommands.

x86: Fault Management For Next Generation AMD Opteron Processors

The fault management feature introduces error-handling and fault-management support for CPUs and memory in systems that use AMD (TM) Opteron and Athlon 64 Rev F processors. These processors are used in the “M2” products from Sun such as the Sun Fire X2200 M2 and Ultra 20 M2. Releases prior to Solaris 10 8/07 provided fault management support for Opteron and Athlon 64 revisions B through E.

Fault management support is enabled by default. The fault management service detects correctable CPU and memory errors, the resulting telemetry is analyzed by diagnosis engines, and errors and faults are corrected whenever possible. When errors cannot be corrected by the system, the extended telemetry provides greater assistance to the system administrator.

For more information see <http://www.opensolaris.org/os/community/fm/>.

x86: Predictive Self-Healing for PCI Express on x64 Systems

Starting with this release, the Solaris OS includes a set of predictive self-healing features to automatically capture and diagnose hardware errors detected on your system.

The Solaris Fault Manager automatically diagnoses failures in x64 hardware. Diagnostic messages are reported by the `fmd` daemon.

For more information about Fault Management in Solaris, see the following:

- [fmd\(1M\)](#) man page
- <http://www.sun.com/msg>
- <http://opensolaris.org/os/community/fm/>

x86: stmsboot Porting

Starting with this release, the `stmsboot` utility is ported to x86 systems. `stmsboot` is a utility that is used to enable or disable MPxIO for fibre-channel devices. This `stmsboot` utility already exists on SPARC systems.

Users can use this utility to enable or disable MPxIO automatically. Previously, users had to enable or disable MPxIO manually, which was difficult, especially for a SAN system boot.

For more information, see the following:

- [stmsboot\(1M\)](#) man page
- Section about “Enabling or Disabling Multipathing on x86 Based Systems” in *Solaris Fibre Channel Storage Configuration and Multipathing Support Guide* at <http://docs.sun.com>.

x86: Concurrent FPDMA READ/WRITE QUEUED Under SATA Module

Starting with this release, concurrent READ/WRITE FPDMA QUEUED commands are supported. There is considerable performance enhancement when performing I/O operations using the Solaris `marvell88sx` driver under specific workload conditions. Other workloads benefit to a smaller degree. There is also significant performance enhancement under many workloads for drives that support this optional portion of the SATA specification.

x86: Tagged Queuing

Tagged queuing enables SATA disks to optimize head motion and performance.

Installation Enhancements

The following installation features and enhancements have been added to the Solaris 10 8/07 release.

NFSv4 Domain Name Configurable During Installation

The NFS version 4 domain can now be defined during the installation of the Solaris OS. In releases prior to Solaris 10 8/07, the NFS domain name was defined during the first system reboot after installation.

The NFSv4 domain name feature affects installation of the OS as follows:

- The `sysidtool` command includes an enhanced `sysidnfs4` program. The `sysidnfs4` program now runs during the installation process to determine whether an NFSv4 domain has been configured for the network.

During an interactive installation, the user is provided with the default NFSv4 domain name that is automatically derived from the OS. The user can accept this default. Or, the user can specify a different NFSv4 domain.

For more information, see the [sysidtool\(1M\)](#) and the [sysidnfs4\(1M\)](#) man pages.

- As part of a Solaris JumpStart™ installation, a new keyword is available in the `sysidcfg` file. The user can now assign a value for the NFSv4 domain by using the new keyword, `nfs4_domain`.

For more information about this new keyword, see the [sysidcfg\(4\)](#) man page. This man page also provides an example of how to use the new `nfs4_domain` keyword.

For more information about the NFSv4 domain name configuration, see the [System Administration Guide: Network Services](#).

Solaris Live Upgrade

Starting with this release, Solaris Live Upgrade has been changed with the following enhancements:

- You can upgrade the Solaris OS when non-global zones are installed on a system by using Solaris Live Upgrade.
- A new package, `SUNWlucfg`, must be installed with the other Solaris Live Upgrade packages `SUNWlur` and `SUNWluu`.

These three packages comprise the software needed to upgrade by using Solaris Live Upgrade. These packages include existing software, new features, and bug fixes. If you do not install these packages on your system before using Solaris Live Upgrade, upgrading to the target release fails.

For more information about upgrading when non-global zones are installed on a system, see [Solaris 10 Installation Guide: Solaris Live Upgrade and Upgrade Planning](#).

Upgrading the Solaris OS When Non-Global Zones Are Installed

Starting with the Solaris 10 8/07 release, you can upgrade the Solaris OS when non-global zones are installed without most of the limitations found in releases prior to Solaris 10 8/07.

Note – The only limitation to upgrading involves a Solaris Flash archive. When you use a Solaris Flash archive to install, an archive that contains non-global zones is not properly installed on your system.

The following changes accommodate systems that have non-global zones installed:

- For the Solaris interactive installation program, you can upgrade or patch a system when non-global zones are installed, with CDs and DVDs. Or you can use a network installation image for either the CDs or DVDs. Previously, you were limited to upgrading with a DVD. The time to upgrade or patch might be extensive, depending on the number of non-global zones that are installed.
- For an automated JumpStart installation, you can upgrade or patch with any keyword that applies to an upgrade or patching. In releases prior to Solaris 10 8/07, only a limited number of keywords could be used. The time to upgrade or patch might be extensive, depending on the number of non-global zones that are installed.
- For Solaris Live Upgrade, you can upgrade or patch a system that contains non-global zones. If you have a system that contains non-global zones, Solaris Live Upgrade is the recommended upgrade program or program to add patches. Other upgrade programs might require extensive upgrade time, because the time required to complete the upgrade increases linearly with the number of installed non-global zones. If you are patching a system with Solaris Live Upgrade, you do not have to take the system to single-user mode and you can maximize your system's uptime.

The following changes accommodate systems that have non-global zones installed:

- A new package, `SUNWlucfg`, must be installed with the other Solaris Live Upgrade packages, `SUNWlur` and `SUNWluu`. This package is required for any system, not just a system with non-global zones installed.

These three packages contain the software needed to upgrade by using Solaris Live Upgrade. These packages include existing software, new features, and bug fixes. If you do not install these packages on your system before using Solaris Live Upgrade, upgrading to the target release fails.

- Creating a new boot environment from the currently running boot environment remains the same with one exception. You can specify a destination disk slice for a shared file system within a non-global zone.

The argument to the `-m` option has a new optional field, `zonename`. The new `zonename` field enables creating the new boot environment and specifying zones that contain separate file systems. This argument places the zone's separate file system on a separate slice in the new boot environment.

- The `lumount` command provides non-global zones with access to their corresponding file systems that exist on inactive boot environments. When the global zone administrator uses the `lumount` command to mount an inactive boot environment, the boot environment is also mounted for non-global zones.
- Listing file systems with the `lufslist` command is enhanced to display a list of file systems for both the global zone and the non-global zones.

A Solaris system that is configured with Trusted Extensions requires extra steps to upgrade labeled zones. For information on this procedure, see “Upgrading a Trusted Extensions System That is Configured with Labeled Zones” under “Installation Enhancements” in *Solaris 10 8/07 Release Notes*.

Keyboard Configuration Automated

Starting with this release, the `sysidkdb` tool configures your USB language and its corresponding keyboard layout.

With the new `sysidkdb` tool, the following procedure occurs:

- If the keyboard is self-identifying, the keyboard language and layout automatically configures during installation.
- If the keyboard is not self-identifying, the `sysidkdb` tool provides you with a list of supported keyboard layouts during installation, so that you can select a layout for keyboard configuration.

Previously, the USB keyboard assumed a self-identifying value of one during the installation. Therefore, all of the keyboards that were not self-identifying always configured for a U.S. English keyboard layout during installation on SPARC.

Note – PS/2 keyboards are not self-identifying. You will have to select the keyboard layout during the installation.

JumpStart Specifications: If the keyboard is not self-identifying and you want to prevent being prompted during your JumpStart installation, select the keyboard language in your `sysidkdb` file. For JumpStart installation, the default is for a U.S. English keyboard layout. To select another language and its corresponding keyboard layout, set the keyboard keyword in your `sysidkdb` file.

For more information, see the [Solaris 10 Installation Guide: Network-Based Installations](#).

Deferred-Activation Patching

Starting with patch 119254-42 and 119255-42, the patch installation utilities, `patchadd` and `patchrm`, have been modified to change the way that certain patches delivering features are handled. This modification affects the installation of these patches on any Solaris 10 release. These “deferred-activation” patches handle the large scope of change delivered in feature patches better.

A limited number of patches are designated as a deferred-activation patch. Typically a deferred-activation patch is a kernel patch associated with a Solaris 10 release after the Solaris 10 3/05 release, such as the Solaris 10 8/07 release. Patches are designated a deferred-activation patch if the variable `SUNW_PATCH_SAFEMODE` is set in the `pkginfo` file. Patches not designated as deferred-activation patches continue to install as before. For example, previously released patches, such as kernel patches 118833-36 (SPARC) and 118855-36 (x86), do not use the deferred-activation patching utilities to install.

Previously, complex patch scripting was required for these kernel patches. The scripting was required to avoid issues during the patch installation process on an active partition because of inconsistencies between the objects the patch delivers and the running system (active partition).

Now, deferred-activation patching uses the loopback file system (`lofs`) to ensure the stability of the running system. When a patch is applied to the running system, the `lofs` preserves stability during the patching process. These large kernel patches have always required a reboot, but now the required reboot activates the changes made by the `lofs`. The patch README provides instructions on which patches require a reboot.

If you are running non-global zones or have `lofs` disabled, consider these points when installing or removing deferred-activation patches:

- All non-global zones must be in a halted state for this patch operation. You must halt the non-global zone before applying the patch.
- Deferred-activation patching requires the loopback file system (`lofs`) in order to complete safely. Systems running Sun Cluster 3.1 or Sun Cluster 3.2 are likely to have `lofs` turned off because of restrictions on HA-NFS functionality when `lofs` is enabled. Therefore, before a deferred-activation patch is installed, you must re-enable the loopback file system by performing the following steps:

1. Remove or comment out the following line in the `/etc/system` file:

```
exclude:lofs.
```

2. Reboot the system.
3. Install the patch.
4. After you have completed the patch installation operation, restore or uncomment the same line from the `/etc/system` file.
5. Reboot the system to resume normal operations.

Note – Sun recommends Solaris Live Upgrade to manage patching. Solaris Live Upgrade prevents the problems of patching a running system. Solaris Live Upgrade reduces the amount of downtime involved in patching and reduces risk by providing fallback capability if problems occur. See *Solaris 10 Installation Guide: Solaris Live Upgrade and Upgrade Planning*.

Networking Enhancements

The following networking features and enhancements have been added to the Solaris 10 8/07 release.

IPsec Tunnel Reform

Solaris now implements IPsec Tunnel Mode per RFC 2401. Inner-packet selectors can be specified on a per-tunnel-interface basis using the new “tunnel” keyword of `ipsecconf(1M)`. IKE and `PF_KEY` handle Tunnel Mode identities for Phase 2/Quick Mode. Interoperability with other IPsec implementations is greatly increased.

For more information, see “[Transport and Tunnel Modes in IPsec](#)” in *System Administration Guide: IP Services*.

Packet Filter Hooks

The packet filter hooks feature includes the following significant functionalities:

- Improved performance in comparison with the STREAMS module approach
- Capability to intercept packets between zones

The packet filter hooks feature is part of a new API that is internal to the kernel. Developers can use the API to work with IP inside the kernel or to intercept packets.

SMF Enhancements to Routing Management

Starting with this release, `routadm(1M)` is enhanced to manage SMF-based routing daemon services. Also, service conversions for the following commands are provided:

- `in.routed(1M)`
- `in.ripngd(1M)`
- `in.rdisc(1M)`
- `in.ndpd(1M)`

As a result, these services can be managed through standard SMF commands such as `svcadm` and `svccfg`, and utilize the restart capabilities that SMF provides.

Quagga Software Routing Suite

Quagga Software Routing Suite delivers a set of IETF routing protocols for Solaris, including OSPF and BGP, allowing for high-availability deployment of Solaris through dynamic routing, manageable through SMF 'routadm'.

Quagga is a community fork of the GNU Zebra software previously included in Solaris, providing many updates and some new features. For more information, see `/etc/quagga/README.Solaris`.

DHCPv6 Client

Starting with this release, the Solaris OS supports Dynamic Host Configuration Protocol for IPv6 (DHCPv6), as described in RFC 3315. DHCPv6 enables Solaris to acquire IPv6 addresses automatically from the local DHCP servers without manual configuration.

For more information, see the following man pages:

- `dhcpage(1M)`
- `in.ndpd(1M)`
- `ifconfig(1M)`
- `ndpd.conf(4)`
- `dhcinfo(1)`

Single Hosts File

Starting with this release, the Solaris OS does not have two separate hosts files. `/etc/inet/hosts` is the single hosts file that contains both IPv4 and IPv6 entries. Solaris system administrators do not need to maintain IPv4 entries in two hosts files that are always synchronized. For backward compatibility, the `/etc/inet/ipnodes` file is replaced with a symbolic link of the same name to `/etc/inet/hosts`.

For more information, see the [hosts\(4\)](#) and the [ipnodes\(4\)](#) man pages.

Large Send Offload

Large Send Offload (LSO) is a hardware off-loading technology. LSO off-loads TCP Segmentation to NIC hardware to improve the network performance by reducing the workload on the CPUs. LSO is helpful for 10Gb network adoption on systems with slow CPU threads or lack of CPU resource. This feature integrates basic LSO framework in Solaris TCP/IP stack, so that any LSO-capable NIC might be enabled with LSO capability.

x86: nge Driver Updated to Support Jumbo Framework

Starting with this release, the nge driver has been updated to enable Jumbo Frame support. The nge driver's default MTU has been raised to 9 Kbytes, that improves system performance and lowers CPU utilization significantly.

For more information, see the [nge\(7D\)](#) man page.

NFSv4 Domain Name Configurable During Installation

For information about this feature, see “[NFSv4 Domain Name Configurable During Installation](#)” on page 86.

Security Enhancements

The following security features and enhancements have been added to the Solaris 10 8/07 release.

Solaris Key Management Framework

The Solaris Key Management Framework (KMF) provides tools and programming interfaces for managing public key (PKI) objects. The `pktool` command enables the administrator to manage PKI objects in `nss`, `pkcs11`, and file-based keystores from a single utility.

The API layer enables the developer to specify the type of keystore to be used. KMF also provides plug-in modules for these PKI technologies. These plug-in modules enable developers to write new applications to use any of the supported keystores.

KMF has a unique feature that provides a system-wide policy database that KMF applications can use regardless of the type of keystore. By using the `kmfcfg` command, the administrator can

create policy definitions in a global database. KMF applications can then choose a policy to enforce, so that all subsequent KMF operations are constrained by the policy being enforced. Policy definitions include rules for the following:

- Strategy for performing validations
- Key usage and extended key usage requirements
- Trust anchor definitions
- OCSP parameters
- CRL DB parameters (for example, location)

For more information, see the following:

- [pktool\(1\)](#) man page
- [kmfcfg\(1\)](#) man page
- Chapter 15, “Solaris Key Management Framework,” in *System Administration Guide: Security Services*

libmd - Message Digest library

Starting with this release, the `libmd` library provides implementations of cryptographic hash algorithms MD4, MD5, SHA1, and SHA2 which comprises SHA256, SHA384, SHA512, by using lightweight APIs. For more information about these APIs and functions offered by `libmd`, see the following man pages:

- [md4\(3EXT\)](#)
- [md5\(3EXT\)](#)
- [sha1\(3EXT\)](#)
- [sha2\(3EXT\)](#)

Solaris Cryptographic Framework

The Solaris Cryptographic Framework feature provides protection of signing keys in a token device. The `elfsign` command also displays more information about signatures and certificates.

For more information, see the [elfsign\(1\)](#) man page.

Solaris Data Encryption Supplement

The Encryption Kit, `SUNWcry` and `SUNWcryr` packages, are included by default with the Solaris 10 8/07 software. Full strength crypto for the Solaris Cryptographic Framework, Kerberos, and OpenSSL is now installed by default.

File System Enhancements

The following file system features and enhancements have been added to the Solaris 10 8/07 release.

Support for iSCSI Target Devices

This Solaris release provides support for iSCSI target devices, which can be disk or tape devices. Releases prior to Solaris 10 8/07 provided support for iSCSI initiators. The advantage of setting up Solaris iSCSI targets is you might have existing fibre-channel devices that can be connected to clients without the cost of fibre-channel HBAs. In addition, systems with dedicated arrays can now export replicated storage with ZFS or UFS file systems.

You can use the `iscsitadm` command to set up and manage your iSCSI target devices. For the disk device that you select as your iSCSI target, you'll need to provide an equivalently sized ZFS or UFS file system as the backing store for the iSCSI daemon.

After the target device is set up, use the `iscsiadm` command to identify your iSCSI targets, which will discover and use the iSCSI target device.

- [iscsiadm\(1M\)](#) man page
- [iscsitadm\(1M\)](#) man page
- [Chapter 14, “Configuring Solaris iSCSI Targets and Initiators \(Tasks\),” in *System Administration Guide: Devices and File Systems*](#)

Extended FILE Space for 32-bit Solaris Processes

The extended FILE space feature supports an addition, F, mode to the `fopen` library command. The F mode enables the opening of files beyond the 255 limit. This feature enables the developers to use the `fopen` command to handle file descriptors up to the limits set using the `limit` or `ulimit` commands.

System Resources Enhancements

The following system resource features and enhancements have been added to the Solaris 10 8/07 release.

1x Branded Zones: Solaris Containers for Linux Applications

Sun's BrandZ technology provides the framework to create non-global branded zones that contain nonnative operating environments. As a simple extension of non-global zones, branded zones offer the same isolated and secure environment, and all brand management is performed through extensions to the current zones structure.

The brand currently available is the 1x brand, Solaris Containers for Linux Applications. These non-global zones provide a Linux application environment on an x86 or x64 machine running the Solaris OS.

The 1x brand includes the tools necessary to install a CentOS 3.5 to 3.8 or Red Hat Enterprise Linux 3.5 to 3.8 inside a non-global zone. Machines running the Solaris OS in either 32-bit or 64-bit mode can execute 32-bit Linux applications.

For more information, see Part III, Branded Zones in the [System Administration Guide: Solaris Containers-Resource Management and Solaris Zones](#).

Also see the following man pages:

- [zoneadm\(1M\)](#)
- [zonecfg\(1M\)](#)
- [brands\(5\)](#)
- [lx\(5\)](#)

Improved zonecfg Procedures for Creating Containers

More integrated resource management and zones features now make it easier to leverage the resource management capabilities of the system through the zonecfg command. The resource configuration you specify is automatically created for you when the zone boots. You no longer have to perform any manual steps related to setting up resource management.

- The zonecfg command can be used to configure resource management settings for the global zone.
- Zone-wide resource controls can be set by using the preferred global property names method. New project and zone resource controls are also available:
 - `zone.max-locked-memory`
 - `zone.max-msg-ids`
 - `zone.max-sem-ids`
 - `zone.max-shm-ids`
 - `zone.max-shm-memory`
 - `zone.max-swap` - Provides swap capping for zones through the capped-memory resource
 - `project.max-locked-memory` - Replaces `project.max-device-locked-memory`
- Some methods have been added for setting the default scheduler in a zone, for example, a new `scheduling-class` property.
- Resource pools have been enhanced. You can add a temporary pool that is created dynamically when a zone boots. The pool is configured through the `dedicated-CPU` resource.
- A `clear` subcommand is available to clear the value for optional settings.
- Enhanced physical memory capping from the global zone is available through improvements to `rcapd(1M)`. Limits are configured through the `capped-memory` resource.

Note – This capability can be used to cap physical memory for lx branded zones and for native zones. For more information, see [“lx Branded Zones: Solaris Containers for Linux Applications”](#) on page 94.

- The resident set size (RSS) accounting has been improved. Improvements have been made to `rcapd`, the resource-capping daemon, and to the `prstat` command.

For more information, see the following:

- [prstat\(1M\)](#) man page
- [rcapd\(1M\)](#) man page
- [zonecfg\(1M\)](#) man page
- [resource_controls\(5\)](#) man page
- *System Administration Guide: Solaris Containers-Resource Management and Solaris Zones*

IP Instances: LAN and VLAN Separation for Non-Global Zones

IP networking can now be configured in two different ways, depending on whether the zone is given is assigned an exclusive IP instance or shares the IP layer configuration and state with the global zone. IP types are configured by using the `zonecfg` command.

The shared-IP type is the default. These zones connect to the same VLANs or same LANs as the global zone and share the IP layer. `lx` branded zones are configured as Shared-IP zones. For more information, see “[lx Branded Zones: Solaris Containers for Linux Applications](#)” on [page 94](#).

Full IP-level functionality is available in an exclusive-IP zone. If a zone must be isolated at the IP layer on the network, then the zone can have an exclusive IP. The exclusive-IP zone can be used to consolidate applications that must communicate on different subnets that are on different VLANs or different LANs.

For more information, see the following:

- [zonecfg\(1M\)](#) man page
- [zones\(5\)](#) man page
- *System Administration Guide: Solaris Containers-Resource Management and Solaris Zones*

For configuration information, see Chapter 17, “Non-Global Zone Configuration (Overview),” in *System Administration Guide: Solaris Containers-Resource Management and Solaris Zones* and Chapter 18, “Planning and Configuring Non-Global Zones (Tasks),” in *System Administration Guide: Solaris Containers-Resource Management and Solaris Zones*.

For information on feature components, see Chapter 26, “Solaris Zones Administration (Overview),” in *System Administration Guide: Solaris Containers-Resource Management and Solaris Zones* and Chapter 27, “Solaris Zones Administration (Tasks),” in *System Administration Guide: Solaris Containers-Resource Management and Solaris Zones*.

Solaris Zones Boot Enhancements

Solaris Zones Boot Enhancements now support boot arguments as part of boot and reboot. The following boot arguments are supported at this time:

- `-m <smf_options>`
- `-i </path/to/init/>`

- -s

Boot arguments can be passed in the following ways:

- `global# zoneadm -z myzone boot -- -m verbose`
- `global# zoneadm -z myzone reboot -- -m verbose`
- `myzone# reboot -- -m verbose`

Boot arguments can also be persistently specified by using the new `bootargs` property in the `zonecfg` command:

```
zonecfg:myzone> set bootargs="-m verbose"
```

This setting will be applied unless overridden by the `reboot`, `zoneadm boot` or `zoneadm reboot` commands.

For more information on boot arguments and the `bootargs` property, see the following:

- [zoneadm\(1M\) man page](#)
- [zonecfg\(1M\) man page](#)
- *System Administration Guide: Solaris Containers-Resource Management and Solaris Zones*

System V Resource Controls for Zones

To limit the total amount of System V resources used by processes within a non-global zone, the following zone-wide resource controls are now included:

- `zone.max-shm-memory`
- `zone.max-shm-ids`
- `zone.max-msg-ids`
- `zone.max-sem-ids`

The resource controls are set through the `add rctl` resource property in `zonecfg` command for non-global zones.

To limit the global zone's consumption, the resource controls can be set through the `prctl` command.

For more information, see the following:

- [prctl\(1\) man page](#)
- [zonecfg\(1M\) man page](#)
- [resource_controls\(5\) man page](#)
- *System Administration Guide: Solaris Containers-Resource Management and Solaris Zones*

Zone Unique Identifier

The Solaris system automatically attaches a globally unique identifier to each non-global zone when the zone is installed. This identifier can be retrieved both in the global zone and in the non-global zone by use of the `zoneadm list -p` command. Users can utilize the zone unique identifier for asset tracking by treating the zone as an asset by itself. This identifier can also be used for identification of zones across the following actions:

- Moving of zones.
- Renaming zones.
- All events that do not involve destruction of zone contents.

For more information, see the [zoneadm\(1M\)](#) man page.

Ability to Mark Zones as “Incomplete”

Starting with this release, users can mark zones as “incomplete” using a new `zoneadm` feature. This new `zoneadm` feature enables the recording of a fatal or permanent zone failure state by administrative software that updates the zone contents.

For more information, see the [zoneadm\(1M\)](#) man page.

Using DTrace in a Non-Global Zone

DTrace can now be used in a non-global zone when the `dttrace_proc` and `dttrace_user` privileges are assigned to the zone. DTrace providers and actions are limited in scope to the zone. With the `dttrace_proc` privilege, `fasttrap` and `pid` providers can be used. With the `dttrace_user` privilege, 'profile' and 'syscall' providers can be used.

You can add these privileges to the set of privileges available in the non-global zone by using the `limitpriv` property of the `zonecfg` command.

“[Configurable Privileges for Non-Global Zones](#)” on page 112 provides an overview of privileges in a non-global zone.

For more information about zone configuration, specifying zone privileges, and using the DTrace utility, see:

- *System Administration Guide: Solaris Containers-Resource Management and Solaris Zones*
- *Solaris Dynamic Tracing Guide*
- [zonecfg\(1M\)](#) man page
- [dttrace\(1M\)](#) man page

Desktop Tools Enhancements

The following desktop tools features and enhancements have been added to the Solaris 10 8/07 release.

Thunderbird 2.0

Thunderbird 2.0 is a full-featured email, RSS, and newsgroup client developed by the Mozilla community. It provides functionality equivalent to the Mozilla mail and newsgroups features.

Firefox 2.0 Web Browser

Firefox 2.0 focuses on user interface innovations that help users in their common browsing tasks while interacting with search, bookmarks and history. Firefox 2.0 has improvements to tabbed browsing, RSS handling, managing extensions, security and performance.

Gaim OTR plugin

Starting with this release a new plug-in, Off-the-Record (OTR), has been added to GAIM.

OTR messaging enables users to have private conversations over GAIM and all its supported messenger services by providing the following:

- Encryption
- Authentication
- Deniability
- Perfect forward secrecy

For more information, see <http://www.cypherpunks.ca/otr/>.

x86: XVideo Support RealPlayer

Starting with this release, XVideo support on RealPlayer improves video playback performance greatly on x86 systems.

X11 Windowing Enhancements

The following X11 windowing features and enhancements have been added to the Solaris 10 8/07 release.

dtlogin Language Selection Overhaul

CDE currently lists the cryptic locale names in the form of a cascade menu in the login screen. The `dtlogin` language selection overhaul feature delivers a more user-friendly language-oriented login list. CDE has a feature to remember the per display default login language name. For SunRay environments, you can use an X resource to disable displays from remembering the login language.

For more information, see the `dtlogin` man page.

X Server DTrace Provider

Starting with this release, the X Window System servers include a User-land Statically Defined Tracing (USDT) DTrace provider for instrumenting X11 client connections. The X Window System servers include the following:

- Xorg
- Xsun
- Xprt
- Xnest
- Xvfb

For more information about available probes and their arguments, and sample DTrace scripts using them, see <http://people.freedesktop.org/~alanc/dtrace/>.

Xorg X11R7.2 Server and Drivers

The Xorg server for the X11 window system, the associated graphics, and the input device drivers, have been upgraded to the X11R7.2 release. The X11R7.2 release includes the Xorg server version 1.2. This release also adds 64-bit versions of the Xorg server for both x64 and SPARC platforms, though drivers for common SPARC graphics devices are not yet available for Xorg.

This release also includes the Xephyr nested X server and the Xorg version of Xvfb, both of which are installed in the `/usr/X11/bin` directory. This version of Xorg no longer supports the Low Bandwidth X (LBX) extension. The use of the X tunneling and compression features of `ssh(1)` is suggested for sites that need X displays across extremely bandwidth-limited network links.

Language Support Enhancements

The following language support features and enhancements have been added to the Solaris 10 8/07 release.

Migration of Existing EMEA, Central and South American Locales to Common Locale Data Repository

The locale data for existing European and Middle East and African (EMEA), Central and South American, and Oceania locales has been migrated to Common Locale Data Repository (CLDR) 1.3. This migration improves locale data quality and ensures consistency of locale data across code sets.

For more information about CLDR, see <http://www.unicode.org/cldr>.

Japanese Font Update

Starting with this release, the Japanese HG font has been updated to compliant with JISX0213:2004.

More Japanese `iconv` Modules for Unicode

Starting with this release, the following two types of codeset conversions between the Unicode and Japanese codesets have been added:

- In conversion from or to eucJP, PCK (SJIS), and ms932, `iconv` now supports UTF-16, UCS-2, UTF-32, UCS-4 and their fixed endian variants, such as UTF-16BE and UTF-16LE, and UTF-8.
- `iconv` now supports codeset name `eucJP-ms` to provide conversion between Japanese EUC and Unicode in the same way as Windows. All Unicode encoding variants mentioned previously, are also supported with `eucJP-ms`.

For more information, see the `iconv_ja(5)` man page.

Input Method Switcher Enhancement and EMEA Keyboard Layout Emulation Support

The input method switcher application, `gnome-im-switcher-applet`, is replaced with a stand-alone GTK+ application, `iiim-panel`. `iiim-panel` now starts and resides on the GNOME panel automatically when you log in to the Java Desktop System (Java DS) in UTF-8 or Asian locales. `iiim-panel` can also run in the Common Desktop Environment (CDE).

IIIMF supports language engines that emulate the EMEA keyboard layout such as French, Polish or Dutch.

For more information, see the online help of the input method preference editor (`iiim-properties`).

x86: Zero-CountryCode Keyboard Layout Support

This feature provides a new command option `kbd -s language`. This option enables users to configure keyboard layouts in kernel. The Zero-CountryCode keyboard layout feature is particularly useful on SPARC systems. In prior releases, all “non-self-ID keyboards” were always recognized as US layout keyboard on SPARC systems.

For more information, see `kbd(1)` man page.

Developer Tool Enhancements

The following developer tools features and enhancements have been added to the Solaris 10 8/07 release.

SunVTS 6.4

SunVTS™ (Sun Validation Test Suite) is a comprehensive software diagnostic package that tests and validates the Sun x86 and SPARC hardware. The SunVTS software verifies the configuration and proper functioning of controllers, devices, and platforms.

Major changes to the Solaris OS for SunVTS include:

- New tests, `xnetlbttest` and `iobustest` have been added. In releases prior to Solaris 10 8/07 both these tests were available only as part of the internal manufacturing package.
- SunVTS memory tests integrated with Test Hang Mitigation (THM) library.

- `nettest` enhancements with a new option to accept packet size.
- `bmcenvironment` test enhancements to support LED tests.
- `netlbttest` has been changed to accommodate `crc` bytes under `nxge` driver.
- `disktest` enhancements
- Generic `tapetest` with improved option settings.
- `iobustest` enhancements which include EFI disc support, bus-related performance counters, stress SIU/NCU, higher stress level cover, PCI-E scan capability.

For more information about these features and tests, see the SunVTS 6.4 documentation at <http://www.sun.com/documentation>.

Driver Enhancements

The following new drivers and driver enhancements have been added to the Solaris 10 8/07 release.

Reliable Datagram Sockets

Starting with this release, Reliable Datagram Sockets (RDS) is a new protocol family that enables a socket to send messages to multiple destinations reliably over the InfiniBand interconnect.

RDS is delivered through a new `SUNWrds` package. The `SUNWrds` package consists of the `rds` and `rdsib` drivers for the socket and transport interface respectively.

Enhanced USB EHCI Host Controller Driver

Enhanced USB EHCI host controller driver provides isochronous transfer support for USB 2.0 or high speed isochronous devices.

For more information, see the `usb_isoc_request(9S)` man page.

USCSI LUN Reset Support

This feature is to supply the logical unit number (LUN) reset support by `uscsi` commands. Users can use LUN reset commands with `uscsi_flags` set as `USCSI_RESET_LUN` with this feature.

SATA HBA Framework and Marvell Driver

Starting with this release, `READ/WRITE FPDMA QUEUED` commands are supported. There is a considerable performance enhancement when performing I/O operations using the Marvell driver under specific work load conditions. Other workloads benefit to a smaller degree. The Sun Branded Hitachi 250GB HDS7225SBSUN250G drives have considerably better performance with larger writes using this feature.

There is also significant performance enhancement under many workloads for drives that support this optional portion of the SATA specification.

Compact Flash Support

The Compact Flash (CF) support feature enables you to use a CF card as an ATA disk through a CF-ATA adapter. This feature helps you start your system from a CF card and store your data on a CF card easily.

For more information about Compact Flash supporting, see the [ata\(7D\)](#) man page.

ACM driver of USB Communication Device Class

Starting with this release, the `usbsacm` driver supports USB modems that conform to the Universal Serial Bus Communication Device Class Abstract Control Model (USB CDC ACM) specification. Customers can attach the `usbsacm` driver with their mobile phones, PCMCIA cards, or any modem-like devices. The `usbsacm` driver outputs term nodes under `/dev/term/`. Customers can then use `pppd(1M)` to transmit datagrams over these serial ports.

CardBus Support

The CardBus support feature adds 32-bit PC Card support in Solaris. Both 16-bit and 32-bit PC Cards are now recognized by Solaris. For more information, see the [pcic\(7D\)](#), [pcic\(7D\)](#) and [cardbus\(4\)](#) man pages.

IBM LTO-4 Tape Drive Support

Starting with this release, the Solaris OS supports the IBM LTO-4 tape drive.

HP LTO-4 Tape Drive Support

Starting with this release, the Solaris OS supports the HP LTO-4 tape drive.

NVIDIA Accelerated Graphics Drivers

Starting with this release, the accelerated graphics drivers for Xorg and OpenGL for NVIDIA Quadro and GeForce cards are included. The `nvidia-settings` and `nvidia-xconfig` configuration tools for these drivers are also provided.

SPARC: `ntwdt` Driver for UltraSPARC-T1 (Niagara) Systems

Starting with release, there is a user-programmable watchdog timer on sun4v platforms that supports backward compatibility. The user can manipulate the Application Watchdog Timer through IOCTLs provided by the backward compatible `ntwdt` pseudo driver.

x86: ACPI Thermal Zone Monitor

The minimal ACPI thermal zone monitor pseudo driver for the Solaris OS handles thermal zone events from the ACPI. Thermal zone events are primarily critical temperature events. If the BIOS in a given system implements specific ACPI methods, this pseudo driver handles the thermal zone events.

x86: Adaptec aac Hardware Support

The updated aac driver supports the new generation, rocket chip-based Adaptec Hardware RAID adapter. The aac driver also supports the Adaptec Storage Management Utility (ASM), which configures and monitors the controller and attached hard drives.

For more information, see the Adaptec web-site <http://www.adaptec.com/en-US/products/adps/>.

x86: Solaris Audio Driver for ATI IXP400

The audioixp driver is the Solaris audio driver for the ATI IXP400 Southbridge chipset from ATI Corporation. The ATI IXP400 chipset includes an embedded AC97 audio controller. This chipset is widely adopted by many motherboard vendors, for example, the Ferrari4000 model. The audioixp driver follows the Solaris Audio Driver Architecture (SADA) frame.

x86: High-Definition Audio Driver

High-definition audio driver, audiohd(7d), is enhanced to support more audio CODECs and provide basic audio playback and recording functionality. The supported high-definition audio CODECs include the following:

- Realtek ALC260/262/880/882/883/885/888
- IDT/Sigmatel STAC9200(D)
- Analog Devices AD1986/1988

x86: SATA AHCI HBA Driver

AHCI is a SATA HBA hot-plug capable driver for SATA controllers that are compatible with the AHCI specification. The AHCI driver supports the INTEL ICH6 and VIA vt8251 controllers, however other AHCI compliant controllers might also work.

For more information, see the [ahci\(7D\)](#) man page.

System Performance Enhancements

The following system performance features and enhancements have been added to the Solaris 10 8/07 release.

SPARC: UltraSPARC T2 PCI Express Interface Unit Performance Counter Data

UltraSPARC T2 systems PCI Express Interface Units (PIU) have built-in performance counters which can be dumped by using `busstat`. The output of the `busstat -l` command shows the following devices for such systems:

- imu#
- mmu#
- peu#

- `bterr#`

where # is an instance number.

The use of this built-in performance counter is intended mainly for Sun field service personnel.

Hashed Cache Index Mode Support

Hashed Cache Index mode is a new hardware feature available in UltraSPARC T2 processors. The hardware uses many more address bits to compute an L2 cache index. As a result, there are more page colors for large pages.

To achieve optimum performance, the Solaris kernel must maximize the number of page colors used by all the threads sharing a cache. The Solaris virtual memory subsystem has been extended to support this new hardware feature. Correct color calculation improves the performance and throughput consistency of application programs on UltraSPARC T2 systems.

Multi-level CMT Scheduling Optimizations

The multi-level Chip Multi-Threaded (CMT) scheduling optimizations feature provides the Solaris kernel with a platform independent mechanism. This mechanism enables discovering and optimizing various performance relevant hardware-sharing relationships existing between CPUs on current and emerging CMT processor architectures, including Niagara II.

This feature also enhances the kernel's thread scheduler or dispatcher with a multilevel CMT load-balancing policy that benefits system performance on various multithreaded, multicore, and multsocket processor-based systems.

For more information on this feature, see the OpenSolaris performance community website, <http://www.opensolaris.org/os/community/performance>.

Process Count Scalability

The process count scalability feature improves the process count scalability of the Solaris OS. Currently, all UltraSPARC systems support a maximum of 8192 contexts. When the number of processes exceeds 8192, then the kernel steals contexts to keep the processes running. Stealing a context from a process involves the following tasks:

- Cross-calling all CPUs that the process ran on
- Invalidating the context for CPUs that are running threads of the process
- Flushing the context from the TLBs of all CPUs that are running threads of the process

This procedure is very expensive and gets worse as the number of processes rise beyond 8K. The process count scalability feature completely redesigns context management. The contexts are managed on a per-MMU basis rather than a global basis which enables efficient TLB flushing and greatly improves the scalability of context management.

The process count scalability feature also improves throughput on workloads that consist of more than 8K active processes, or create and destroy processes at a high rate, and is most beneficial on systems with many CPUs.

MPSS Extended to Shared Memory

The multiple page size support (MPSS) for shared memory feature adds large page support for mapping shared memory and provides an out-of-box (OOB) policy for the use of large pages for shared memory. The MPSS support is for shared memory created by the `mmap(1)` of `/dev/zero` or with the `MAP_ANON` flag, and for System V shared memory. This feature also adds support for `mencnt l(2)` changing the page size of these shared memory segments.

MPSS support is also extended for the use of large pages for memory created by the `mmap(1)`, `mmap(MAP_PRIVATE)` of `/dev/zero`.

Device Management Enhancements

The following device management features and enhancements have been added to the Solaris 10 8/07 release.

Enhanced `st` SCSI Reservations

Starting with this release, there is a new reservation mechanism in the `st` driver. The new mechanism enables the `st` driver to reserve the tape drive only when a command that requires reservations is sent. The reservation mechanism also enables the `st` driver to process inquiry commands issued from other hosts while the drive is reserved by a different host.

Some of Independent Software Vendor's (ISV's) backup software and media management tools benefit from the enhanced `st` SCSI reservations feature. Because of this new feature, management tools could inquire and browse tape libraries when the backup tool is reading or writing tapes.

CPU Power Management

This feature introduces two new `power.conf` keywords to permit CPU devices to be power managed independently of automatic power management. These are the new `power.conf` keywords:

- `cpupm`

Usage:

```
cpupm <behavior>
```

Here, the behavior is `enable` or `disable`.

For backward compatibility, if the `cpupm` keyword is not present in the `/etc/power.conf` file, the CPUs are power managed if `autopm` is enabled, and not power managed if `autopm` is disabled. `enable` or `disable` are independent of the `autopm` setting.

- `cpu-threshold`

Usage:

```
cpu-threshold <threshold>
```

This keyword enables the user to specify a threshold which will apply to any power-manageable CPU, independent of the `system-threshold` value.

If CPU power management is enabled, the power level of any CPU that is idle for the specified threshold time is reduced to the next lower level.

If `cpu-threshold` is absent, system threshold is used.

For more information, see the [power.conf\(4\)](#) man page.

Console Subsystem Enhancements

The following console subsystem enhancement has been added to the Solaris Solaris 10 8/07 release.

Coherent Console

The coherent console feature implements a part of the kernel console subsystem to facilitate rendering console output. The coherent console uses the Solaris kernel mechanisms to render console output rather than Programmable Read-Only Memory (PROM) interfaces. This reduces the console rendering dependence on OnBoot PROM (OBP).

The coherent console uses a kernel-resident `framebuffer` driver to generate console output. The generated console output is more efficient than using OBP rendering. The coherent console also avoids idling CPUs during SPARC console output and enhances the user experience.

For example, the coherent console increases the SPARC console text throughput and scrolling rate and provides ANSI color.

What's New in the Solaris 10 11/06 Release

This chapter summarizes all the new features in the Solaris 10 11/06 release.

System Administration Enhancements

The following system administration features and enhancements have been added to the Solaris 10 11/06 release.

Storage Networking Industry Association Multipath Management API Support

This feature provides Sun's implementation of the Storage Networking Industry Association (SNIA) Multipath Management API (MP API). The support includes the following:

- MP API common library
- Plug-in library for Solaris native multipathing solution - MPxIO/scsi_vhci driver
- mpathadm CLI

The MP API common library exports the defined set of standard interfaces. The plug-in library for the scsi_vhci driver enables scsi_vhci multipathing devices to be administered through MP API and its associated CLI, `mpathadm`.

The SNIA MP API defines standard interfaces for multipathing discovery and administration which enables multipathing management applications to use the common set of APIs across vendor-unique multipathing solutions on Solaris. Sun is supplying a plug-in library that enables Solaris native multipathing solution to be managed through the API and associated CLI.

Sun Java Web Console Changes

The Sun Java™ Web Console provides a common location for users to work with web-based management applications. Users access the console by logging in through an HTTPS port, using one of several supported web browsers. The single entry point that is provided by the console eliminates having to learn URLs for multiple applications. The console provides authentication and authorization services for all applications that are registered with the console.

All console-based applications conform to the same user interface guidelines. The Sun Java Web Console also provides auditing and logging services for all registered applications.

The Solaris ZFS Administration tool is a console application that is provided beginning with the Solaris 10 6/06 release. For more information about using the Solaris ZFS web-based management tool, see *Solaris ZFS Administration Guide*.

Starting with the Solaris 10 11/06 release, the Sun Java Web Console includes the following changes:

- The console now supports applications that are based on the JavaServer™ Faces technology.
- The console server is configured to run as a service that is managed by the Service Management Facility (SMF). SMF commands can now be used to manage the console web server by using the Fault Managed Resource Identifier (FMRI) “system/webconsole:console.” The `smcwebserver` command can also be used to start, stop, enable, and disable the console server, as in previous Solaris 10 releases.

For more information, see the `smcwebserver(1M)` man page.

- A new command, `wcadmin`, is used to configure console properties. The command is also used to deploy and enable console applications that are written for the new version of the console. The `smreg` command, which was previously used to perform similar tasks, is now used only to register and unregister the applications that were developed for previous versions of the console.

For more information, see the `smreg(1M)` and `wcadmin(1M)` man pages.

For more information, see “Working With the Sun Java Web Console (Tasks)” in *System Administration Guide: Basic Administration*.

File-System Monitoring Tool

This file system enhancement is new in the Solaris 10 11/06 release.

A new file-system monitoring tool, `fsstat`, is available to report file-system operations. Activity can be reported by mount point or by file-system type.

For more information, see the `fsstat(1M)` man page.

System Resource Enhancements

The following system resource features and enhancements have been added to the Solaris 10 11/06 release.

Resource Management Features

The following resource management features and enhancements have been added to the Solaris 10 11/06 release.

Resource Pools Facility Service FMRI

Resource pools and dynamic resource pools have been integrated into the Solaris service management facility (SMF). Dynamic resource pools are now enabled separately of the resource pools service.

The dynamic resource pools service fault management resource identifier (FMRI) is `svc:/system/pools/dynamic`. The resource pools service FMRI is `svc:/system/pools`.

The enabling and disabling mechanisms through `pooladm(1M)` are still available.

Note – When a system is upgraded, if a `/etc/pooladm.conf` file exists, the configuration contained in the file is applied to the system.

For more information, see:

- [System Administration Guide: Solaris Containers-Resource Management and Solaris Zones](#)
- `pooladm(1M)` man page
- `poold(1M)` man page
- `libpool(3LIB)` man page
- `smf(5)` man page

Solaris Zones Features

The following Solaris zones features and enhancements have been added to the Solaris 10 11/06 release.

Solaris Zones Renaming Feature

The zone name is now an attribute that can be set through the `zonecfg` command. Only zones in the configured or installed states can be renamed.

For information about zones configuration and zone states, see:

- [System Administration Guide: Solaris Containers-Resource Management and Solaris Zones](#)
- `zonecfg(1M)` man page
- `zones(5)` man page

Zones Move and Clone Features

Two new subcommands, `move` and `clone`, have been added to the `zoneadm` command. You can now do the following:

- Relocate a non-global zone from one point on a system to another point on the same system
- Rapidly provision a new non-global zone based on the configuration of an existing zone on the same system

For more information, see:

- [System Administration Guide: Solaris Containers-Resource Management and Solaris Zones](#)
- `zoneadm(1M)` man page

Migrating a Non-Global Zone From One Machine to Another

The `zonecfg` and `zoneadm` commands have been modified to enable you to migrate a non-global zone from one system to another. The procedure used detaches a halted zone from its current location, and attaches the zone at a new location. The global zone on the target system must be running the following:

- The same release as the original host
- The same versions of operating system packages and patches as the original host

The zone detach process creates the information necessary to attach the zone on a different system. The zone attach process verifies that the new machine has the correct configuration to host the zone. You can make the `zonpath` available on the new host in several ways. Therefore, the actual movement of the `zonpath` from one system to another is a manual process that is performed by the zone administrator.

When attached to the new system, the zone is in the installed state.

For more information, see:

- [System Administration Guide: Solaris Containers-Resource Management and Solaris Zones](#)
- `zonecfg(1M)` man page
- `zoneadm(1M)` man page

Configurable Privileges for Non-Global Zones

The `limitpriv` property of the `zonecfg` command can be used to specify the set of privileges that processes are limited to in a non-global zone.

You can do the following:

- Augment the default set of privileges with the understanding that such changes might allow processes in one zone to affect processes in other zones by being able to control a global resource
- Create a zone with fewer privileges than the default, safe set

For more information about configuring privileges for zones and zone privilege restrictions, see:

- [System Administration Guide: Solaris Containers-Resource Management and Solaris Zones](#)
- `zonecfg(1M)` man page

Note – Note the following:

- Non-global zones are still booted with the standard set of safe privileges by default.
 - One set of privileges cannot be removed from the zone's privilege set, and another set of privileges cannot be included in the zone's privilege set
-

Logical Domains Features

The following logical domains features and enhancements have been added to the Solaris 10 11/06 release.

Logical Domains (LDDoms) 1.0 Software

The Logical Domains (LDDoms) 1.0 software enables system administrators to create and manage logical domains. This software provides multiple software partition support and the following features for Sun4v-based platforms:

- Software upgrade to UltraSPARC T1 systems (Solaris 10 11/06 and firmware upgrade)
- Up to 32 logical domains per system, managed by a CLI, the Logical Domains (LDDoms) Manager 1.0 software, which is a separate download
- Each guest domain can be created, destroyed, reconfigured, and rebooted independently
- Virtual console, Ethernet, disk, and cryptographic acceleration
- Live dynamic reconfiguration of virtual CPUs
- Fault management architecture (FMA) diagnosis for each logical domain

In addition to the Solaris 10 11/06 OS, a minimum level of system firmware 6.4 and Logical Domains Manager 1.0 software is required to have Logical Domains functionality.

Security Enhancements

The following security features and enhancements have been added to the Solaris 10 11/06 release.

Solaris Trusted Extensions

The Solaris Trusted Extensions software provides multilevel security for the Solaris OS, including mandatory access control for the following:

- Files
- File systems
- Processes
- Removable devices
- Networking
- Desktop environments

- Printing

The Solaris Trusted Extensions software also provides tools for the following actions:

- Defining policies
- Setting up sensitivity labels
- Performing trusted system management

The Solaris Trusted Extensions feature enables you to define your data access policies to control information in a flexible but highly secure manner. Solaris Trusted Extensions can be used as a configuration option for the Solaris OS.

For more information about Solaris Trusted Extensions, <http://www.sun.com/smi/Press/sunflash/2006-02/sunflash.20060214.3.xml>.

Solaris Trusted Extensions for Printing

The Solaris Trusted Extensions for printing feature enables the following features:

- Restricted output to printers by label ranges
- Specially labeled banner and trailer pages
- Specially labeled headers and footers

Solaris Trusted Extensions File-System Labeling

Starting with this release, files and directories are labeled by the zone or host that exports them. The mount policy is restricted to prevent writing down.

Device Management Enhancements

The following device management features and enhancements have been added to the Solaris 10 11/06 release.

Support for PCI Express (PCIe)

This Solaris release provides support for the PCI Express (PCIe) interconnect for both SPARC and x86 systems.

PCIe is designed to connect peripheral devices to desktop, enterprise, mobile, communication, and embedded applications.

The PCIe interconnect is an industry-standard, high-performance, serial I/O bus.

The PCIe software provides the following features in this Solaris release:

- Support for extended PCIe configuration space
- Support for PCIe baseline error handling and MSI interrupts
- Modified IEEE-1275 properties for PCIe devices

- PCIe hot-plug support (both native and ACPI-based) by enhancing the `cfgadm_pci` component of the `cfgadm` command
- ATTN button usage-based PCIe peripheral autoconfiguration

The following `cfgadm` example output displays the hot-pluggable PCIe devices on x86 systems. Note that the following display might differ from platform to platform. Check your hardware platform guide for the correct `cfgadm` syntax.

```
# cfgadm pci
Ap_Id                Type      Receptacle  Occupant    Condition
pcie1                unknown  empty       unconfigured unknown
pcie2                unknown  empty       unconfigured unknown
pcie3                unknown  empty       unconfigured unknown
pcie4                etherne/hp  connected  configured  ok
pcie5                pci-pci/hp  connected  configured  ok
pcie6                unknown  disconnected unconfigured unknown
```

The administrative model for hot-plugging PCIe peripherals is the same as for PCI peripherals, which use the `cfgadm` command.

For more information, see the `cfgadm_pci(1M)` man page and *System Administration Guide: Devices and File Systems*. Check your hardware platform guide to ensure that PCIe and PCIe hot-plug support is provided on your system. In addition, carefully review the instructions for physically inserting or removing adapters on your system and the semantics of device autoconfiguration, if applicable.

For more information about PCIe technology, see <http://www.pcisig.com>.

x86: Sun Fire X4500 SATA Disk FMA

A new Fault Management Architecture-based diagnosis engine (DE) is provided on the Sun Fire X4500. This DE monitors the disk drives for predictive failures by using the SMART technology in the disk drive's own firmware. When a disk failure is imminent, the LED next to the disk is illuminated and a Fault Management Architecture fault is generated. This fault alerts the administrator to take specific action to ensure system availability and full performance.

SPARC: Transitioning SPARC-Based Systems From `ipge` to `E1000g` Network Drivers

`ipge` drivers are used in Ontario and other SPARC platforms that have the NorthStar card installed. `E1000g` drivers are used in all other platforms.

Starting with this release, Ontario and other SPARC based platforms transition from `ipge` to `e1000g` drivers. This feature makes the `e1000g` the default driver for all Sun platforms that use Intel 1G chipsets. With the transition, the customer does not need to know which platform the `ipge` or the `e1000g` driver covers or which driver to install in a particular platform. This feature reduces system management complexity.

For more information, see “*Certain 3rd Party Applications May Break on Transition From `ipge` to `e1000g` Network Driver*” on <http://sunsolve.sun.com/>.

Solaris Fibre Channel Host-Based Logical Unit Number Masking

The Solaris fibre channel logical unit number (LUN) masking feature enables system administrators to prevent the kernel from creating device nodes for specific unapproved LUNs.

For more information, see the [fp\(7d\)](#) man page.

SPARC: Extended Message Signaled Interrupt Support for Fire-Based Platforms

Extended Message Signaled Interrupts (MSI-X) are an enhanced version of MSI interrupts. With MSI-X support, device driver writers have a choice between MSI and MSI-X interrupts. MSI-X interrupts are now supported on SPARC PCI-Express platforms (Ultra 45 and Sun Fire T2000). The Sun Fire T2000 might also include the Sun Fire T1000 machine.

A new `mdb/kmdb` debugger command, `::interrupts`, is also provided to retrieve a device's registered interrupt information on supported SPARC and x86 systems.

For more information, see “Interrupt Handlers” in *Writing Device Drivers*.

Improved Device in Use Error Checking

The following utilities have been enhanced to detect when a specified device is in use:

- `dumpadm`
- `format`
- `mkfs` and `newfs`
- `swap`

These enhancements mean that these utilities might detect some of the following usage scenarios:

- Device is part of a ZFS storage pool
- Device is a dump or swap device
- Mounted file system or an entry for the device exists in the `/etc/vfstab` file
- Device is part of a live upgrade configuration
- Device is part of a Solaris Volume Manager configuration or Veritas Volume Manager configuration

For example, if you attempt to use the `format` utility to access an active device, you will see a message similar to the following:

```
# format
.
.
.
Specify disk (enter its number): 1
selecting c0t1d0
```

```
[disk formatted]
Warning: Current Disk has mounted partitions.
/dev/dsk/c0t1d0s0 is currently mounted on /. Please see umount(1M).
/dev/dsk/c0t1d0s1 is currently used by swap. Please see swap(1M).
```

However, these utilities do not detect all scenarios in the same way. For example, you can use the `newfs` command to create a new file system on a device in a live upgrade configuration. You cannot use the `newfs` command to create a new file system on a device that is part of a live upgrade configuration if it also has a mounted file system.

Desktop Enhancements

The following desktop features and enhancements have been added to the Solaris 10 11/06 release.

Default Desktop Session in `dtlogin`

Starting with this release, when a user logs in to the Solaris Desktop for the first time, Sun Java Desktop System (Java DS) is the default desktop environment instead of the Common Desktop Environment (CDE). Java DS has also become the default environment for users who chose a desktop environment on an earlier Solaris release that is no longer present in this Solaris release, such as OpenWindows™ or GNOME 2.0.

System administrators can modify the `dtlogin` configuration to override the default choices by using the `defaultDt` and `fallbackDt` resources.

For more information about `defaultDt` and `fallbackDt` resources, see the `dtlogin(1M)` man page.

Adobe Flash Player Plugin for Solaris

The Adobe Flash Player, formerly known as Macromedia Flash Player is the standard for delivering high-impact and rich web content. Designs, animation, and application user interfaces are deployed immediately across all browsers and platforms, attracting and engaging users with a rich web experience.

GNOME-VFS and Nautilus ACL Support

Starting with this release, ACL support has been added to GNOME-VFS and Nautilus. The GNOME file manager now enables the file system access control lists to be accessed and modified. The GNOME-VFS and Nautilus ACL support feature brings an existing file system functionality to the desktop.

Solaris Trusted Extensions Desktops

Labeled security has been extended to the two desktop interfaces in the Solaris 10 11/06 release. Users have access to both Trusted Java Desktop System (Trusted Java DS) and Trusted Common Desktop Environment (Trusted CDE) which include the following features:

- Multilevel sessions to enable users to access the data that they are authorized to see without compromising security
- Trusted path verification for to assure that the user session is not being hijacked
- Labeled windows to display the label of a window or a document
- Drag-and-drop security enforcement to ensure data movement is controlled and users are informed about security violations
- Labeled device allocation for CD-ROM, DVD, audio, and other devices to limit the transfer of sensitive data to insecure devices
- Secured remote access to multilevel sessions and single level sessions from other systems

Installation Enhancements

The following installation features and enhancements have been added to the Solaris 10 11/06 release.

Solaris Flash Archives

This Solaris Flash enhancement enables a user to create an archive that includes large files. The `flarcreate` command creates a Solaris Flash archive that can contain individual files of 4 Gbytes or more. The available archive utilities are as follows:

- The `cpio` archive utility is the default. Individual files cannot be larger than 2 or 4 Gbytes. The size limitation depends on the version of `cpio` used.
- The portable archive interchange utility, `pax`, is started with the `-L pax` option. If the `-L pax` option is specified, the archive can be created without size limitations on individual files. The `pax` utility was included in the Solaris 7 OS release. The Solaris Flash archives created by using the `pax` utility can only be deployed on a Solaris OS with a `pax` utility. When a user deploys the archive on systems that are running the Solaris 2.6 or earlier versions, the user must use the `cpio` option.

For more information, see the [pax\(1\)](#) and the [cpio\(1\)](#) man pages. See also the *Solaris 10 Installation Guide: Solaris Flash Archives (Creation and Installation)*.

Secure By Default Network Profile

Starting with this release, you can, during installation, set the default behavior for network services to run in a much more secured manner. During an interactive installation (hands on), this security option is provided in the installation configuration selection screens. For automated JumpStart installations (hands off), you can select a limited network profile by using a new `service_profile` keyword in the `sysidcfg` file.

If you choose to restrict network security during the initial installation, numerous services are fully disabled during the installation. Other services are still enabled, but these services are limited to local connections only. Solaris Secure Shell remains available for remote administrative access to the system.

With this limited networking profile, you reduce your risk of exposure on the Internet or LAN. The system retains full graphical desktop use and outbound network access. For example, you can still access your graphical interface, use browsers or email clients, and mount NFSv4 file shares.

The existing service configuration is not altered by an upgrade.

The network services can be easily reopened after installation by using the `net services open` or by enabling individual services by using SMF commands.

For more information about this new security option, see the following references.

TABLE 6-1 Additional Security Information

Administer security for network services	“How to Create an SMF Profile” in <i>System Administration Guide: Basic Administration</i>
Reopen network services after installation	“Revising Security Settings After Installation” in <i>Solaris 10 11/06 Installation Guide: Planning for Installation and Upgrade</i>
Plan installation configuration	“Planning Network Security” in <i>Solaris 10 11/06 Installation Guide: Planning for Installation and Upgrade</i>
Select restricted network security during a hands-on installation	Chapter 2, “Installing With the Solaris Installation Program (Tasks),” in <i>Solaris 10 Installation Guide: Basic Installations</i>
Set up restricted network security for a JumpStart installations	“service_profile Keyword” in <i>Solaris 10 11/06 Installation Guide: Network-Based Installations</i>

Installing Solaris Trusted Extensions

Solaris Trusted Extensions provides multilevel security for the Solaris OS. This feature enables you to control information in a flexible but highly secure manner. You can now enforce strict access controls to your data, based on data sensitivity, not just data ownership.

An installation that accesses Solaris Trusted Extensions differs from a standard installation. For a list of these installation differences and further information about Solaris Trusted Extensions, see [Chapter 3, “Installing Solaris Trusted Extensions Software \(Tasks\),” in *Solaris Trusted Extensions Installation and Configuration for Solaris 10 11/06 and Solaris 10 8/07 Releases*](#).

For more information about Solaris Trusted Extensions, see the README file in the `Solaris_10/ExtraValue/CoBundled/Trusted_Extensions` directory. See also, [“Solaris Trusted Extensions” on page 113](#).

System Performance Enhancements

The following system performance features and enhancements have been added to the Solaris 10 11/06 release.

SPARC: Watchdog Timer for Sun4V

This feature provides system wide watchdog timer functionality. The watchdog timer will be continually reset by the kernel. Failure to reset the timer by the kernel before it expires will result in the system getting reset.

Networking Enhancements

The following networking features and enhancements have been added to the Solaris 10 11/06 release.

Sun Java System Message Queue 3.7 Update 1

Message Queue (MQ) 3.7 Update 1 is a maintenance release for MQ 3.6. This release contains bug fixes as well as performance improvements which reduce disk write overhead for transacted messages.

New and Updated Drivers

The following drivers have been added or enhanced in the Solaris 10 11/06 release.

ST Driver Support for Quantum LTO-2 and LTO-3 Tape Drives

Starting with this release, there is ST driver support for Quantum LTO-2 and LTO-3 Tape Drives.

For more information about the ST driver, see the `st` man page.

CDB Length Capability

HBA drivers can enable target drivers to query the maximum supported CDB length by using `scsi_ifgetcap`. The target driver asks for the capability at attach time and if the HBA driver supports the capability, it returns the maximum length of the CDB in bytes. The target driver can then use this value to make decisions about which CDBs to use for that HBA.

Language Support

The following language support features and enhancements have been added to the Solaris 10 11/06 release.

IIIMF and Language Engines

The Internet Intranet Input Method Framework (IIIMF) has been upgraded from rev.10 to rev.12.

This framework provides the following new features:

- **Input Method Switcher** - This feature displays input method status and switches input languages. You can add the input method switcher to the Java Desktop System (Java DS) panel. Select **Add to Panel -> Utility -> InputMethod Switcher** to add the input method switcher to the Java DS panel.
- **Utility for iiim-properties** - This feature supports various input method preferences. Use one of the following methods to start the `iiim-properties` utility:
 - Select **Launch -> Preferences -> Desktop Preferences -> Input Methods**.
 - Click mouse button 3 on the Input method switcher and select Preference.
 - In the CDE environment, select **Tool -> Input Method Preference** from the CDE main menu or type **`iiim-properties`** at the command prompt.

Each language engine has also been upgraded to the IIIMF rev.12 base. The Japanese language engines, ATOK12 and Wnn6, have been updated to “ATOK for Solaris” and Wnn8 respectively. “ATOK for Solaris” is equivalent to ATOK17. A new Chinese chewing input method has also been added to the IIIMF.

What's New in the Solaris 10 6/06 Release

This chapter summarizes all the new features in the Solaris 10 6/06 release.

System Administration Enhancements

The following system administration features and enhancements have been added to the Solaris 10 6/06 release.

The Solaris ZFS File System

This file system enhancement is new in the Solaris 10 6/06 release.

This Solaris update release includes Solaris ZFS, a new 128-bit file system. Solaris ZFS provides simple administration, transactional semantics, end-to-end data integrity, and immense scalability. Solaris ZFS is not an incremental improvement to existing technology. Rather, Solaris ZFS is a fundamentally new approach to data management.

Solaris ZFS uses a pooled-storage model that completely eliminates the concept of volumes. Thus, Solaris ZFS eliminates the associated problems of partition management, provisioning, and growing file systems. Thousands of file systems can all draw from a common storage pool. Each system consumes only as much space as actually needed. The combined I/O bandwidth of all devices in the pool is available to all file systems at all times.

All operations are “copy-on-write” transactions, so the on-disk state is always valid. Every block has a checksum, so silent data corruption is impossible. In addition, the data is self-healing in replicated configurations. This feature means that if one copy is damaged, Solaris ZFS detects the damage and uses another copy to repair the damaged copy.

Solaris ZFS Easy Administration

For system administrators, the greatest improvement of Solaris ZFS over traditional file systems is the ease of administration.

Solaris ZFS takes a single command to set up a mirrored storage pool and file system. For example:

```
# zpool create home mirror c0t1d0 c1t2d0
```

The preceding command creates a mirrored storage pool named `home` and a single file system named `home`. The file system is mounted at `/home`.

With Solaris ZFS, you can use whole disks instead of partitions to create the storage pool.

You can use the `/home` file system hierarchy to create any number of file systems beneath `/home`. For example:

```
# zfs create home/user1
```

For more information, see the [zpool\(1M\)](#) and [zfs\(1M\)](#) man pages.

In addition, Solaris ZFS provides the following administration features:

- Backup and restore capabilities
- Device management support
- Persistent snapshots and cloning features
- Quotas that can be set for file systems
- RBAC-based access control
- Storage pool space reservations for file systems
- Support for Solaris systems that have zones installed

For more information, see the *Solaris ZFS Administration Guide*.

ZFS Web-Based Management

The Solaris 10 6/06 release includes the Solaris ZFS web-based management tool, which enables you to perform much of the administration that you can do with the ZFS command line interface. You can perform the following administrative tasks with the Solaris ZFS administration console:

- Create a new storage pool
- Add capacity to an existing pool
- Move (export) a storage pool to another system
- Import a previously exported storage pool to make the pool available on another system
- View information about storage pools
- Create a file system
- Create a volume
- Take a snapshot of a file system or a volume
- Roll back a file system to a previous snapshot

You can access the Solaris ZFS administration console through a secure web browser at the following URL:

```
https://system-name:6789
```

If you type the appropriate URL and are unable to reach Solaris ZFS administration console, the server might not be started. To start the server, run the following command:

```
# /usr/sbin/smcwebserver start
```

If you want the server to run automatically when the system boots, run the following command:

```
# /usr/sbin/smcwebserver enable
```

Note – You cannot use the Solaris Management Console (smc) to manage ZFS storage pools or file systems.

ZFS and Solaris Zones

The Solaris Zones partitioning technology supports Solaris ZFS components, such as adding Solaris ZFS file systems and storage pools into a zone.

For example, the file system resource type in the `zonecfg` command has been enhanced as follows:

```
zonecfg:myzone> add fs
zonecfg:myzone:fs> set type=zfs
zonecfg:myzone:fs> set dir=/export/share
zonecfg:myzone:fs> set special=tank/home
zonecfg:myzone:fs> end
```

For more information, see the [zonecfg\(1M\)](#) man page and the [Solaris ZFS Administration Guide](#).

Solaris Installation Tool Support of ZFS File Systems

This release supports the following Solaris installation tools:

- **Solaris custom Jumpstart™** - You cannot include Solaris ZFS file systems in a Jumpstart profile. However, you can run the following scripts from a Solaris ZFS storage pool to set up an install server or an install client:
 - `setup_install_server`
 - `add_install_server`
 - `add_install_client`
- **Solaris Live Upgrade** - Preserves your original boot environment and carries over your Solaris ZFS storage pools into the new environment. Currently, Solaris ZFS cannot be used as a bootable root file system. Therefore, your existing Solaris ZFS file systems are not copied into the boot environment (BE).

- **Solaris Initial Install** - Solaris ZFS file systems are not recognized during an initial installation. However, if you do not specify any of the disk devices that contain Solaris ZFS storage pools to be used for the installation, you should be able to recover your storage pools by using the `zpool import` command after the installation. For more information, see the [zpool\(1M\)](#) man page.

As with most reinstallation scenarios, you should back up your Solaris ZFS files before proceeding with the initial installation option.

- **Solaris Upgrade** – Your Solaris ZFS file systems and storage pools are preserved.

New Solaris ACL Model

Solaris ZFS implements a new ACL model. Previous versions of the Solaris OS only supported an ACL model that was primarily based on the POSIX ACL draft specification. The POSIX-draft based ACLs are used to protect UFS files. A new model that is based on the NFSv4 specification is used to protect Solaris ZFS files.

The new ACL model :

- Is based on the NFSv4 specification and the new ACLs that are similar to NT-style ACLs
- Provides a more granular set of access privileges
- Uses the `chmod` and `ls` commands rather than the `set facl` and `get facl` commands to set and display ACLs
- Provides richer inheritance semantics for designating how access privileges are applied from directory to subdirectories, and so on

The recently revised `chmod(1)` man page adds many new examples that demonstrate usage with Solaris ZFS. The `acl(5)` man page has an overview of the new ACL model. In addition, the [Solaris ZFS Administration Guide](#) provides extensive examples of using ACLs to protect ZFS files.

Predictive Self-Healing for x64 Systems

This system administration feature is new in the Solaris 10 6/06 release.

Starting with this release, the Solaris OS includes a set of predictive self-healing features to automatically capture, diagnose, and respond to hardware errors detected on your system.

The Solaris Fault Manager now provides support for CPU and Memory errors detected on x64 systems, including:

- AMD Athlon 64 and Opteron™ CPU errors
- Northbridge and Hypertransport links errors
- DRAM correctable, uncorrectable, and ChipKill errors

The Solaris Fault Manager automatically diagnoses failures in x64 hardware. The Solaris Fault Manager also attempts to automatically offline or isolate a faulty CPU, cache, or DRAM memory region. Diagnostic messages are reported by the `fmd` daemon.

For more information about Fault Management in Solaris, see:

- [fmd\(1M\)](#) man page
- <http://www.sun.com/msg>
- <http://opensolaris.org/os/community/fm/>

Predictive Self-Healing Support for SNMP Notification

This system administration feature is new in the Solaris 10 6/06 release.

Starting with this release, the Solaris OS includes a set of predictive self-healing features to automatically capture, diagnose, and respond to hardware errors detected on your system. The self-healing diagnosis results are reported to the `syslogd` service.

The Solaris Fault Manager, `fmd`, now enables you to do the following:

- Publish diagnostic results through the Solaris System Management Agent (SMA), including SNMP traps
- Search an SNMP MIB for fault management information about each machine

The Fault Management MIB is located at `/etc/sma/snmp/mibs/SUN-FM-MIB.mib` on the Solaris system.

For more information about configuring SNMP on Solaris, see:

- [fmd\(1M\)](#) man page
- [syslogd\(1M\)](#) man page
- *Solaris System Management Agent Administration Guide*
- <http://www.sun.com/msg>

SunVTS 6.2

This system administration feature is new in the Solaris 10 6/06 release.

SunVTS (Sun Validation Test Suite) 6.2 is a comprehensive software diagnostic package that tests and validates Sun SPARC and x86 hardware. This feature verifies the configuration and proper functioning of controllers, devices, and platforms to validate the hardware.

SunVTS supports the following new hardware:

- Sun Crypto Accelerator 6000 board tested by `cryptotest` cryptographic test. The `cryptotest` test is now supported on x86 platforms.
- The UltraSPARC T1 processor tested by the data translation look-aside buffer test (`dtlbtest`) and the RAM test (`ramtest`).

If executed on unsupported hardware, SunVTS stops execution and provides appropriate warning messages. This feature is an x86 enhancement for SunVTS.

For more SunVTS 6.2 documentation, see <http://www.sun.com/documentation>.

Common Agent Container

This system administration feature is new in the Solaris 10 6/06 release.

The Common Agent Container (CAC) is a stand-alone Java™ program that implements a container for Java management applications. CAC provides a management infrastructure designed for the management functionality based on Java Management Extensions (JMX™) and Java Dynamic Management Kit (JDMK). The `SUNCacao` package installs the CAC software in the `/usr/lib/cacao` directory. Typically, CAC is not visible to the user or administrator.

Two occasions when an administrator might need to interact with the container daemon are:

- When an application tries to use a network port that is reserved for the CAC
- To regenerate the CAC certificate keys when a certificate store is compromised

For more information, see the [Chapter 14, “Troubleshooting Software Problems \(Overview\)”](#) in *System Administration Guide: Advanced Administration*.

iSCSI Logout Support

This system administration feature is new in the Solaris 10 6/06 release.

The iSCSI logout support feature allows a user to logout from an iSCSI target without rebooting the host. When a user tries to remove or disable a discovery method or address and the target is not in use, then the target logs out and cleans up all related resources. If the target is in use, then the discovery address or method remains enabled and the *logical unit in use* message is logged. This feature introduces a new behavior to safely log out of unused devices without rebooting the host.

The following commands can be use to apply this feature:

- `iscsiadm modify discovery -[tsi] disable`
- `iscsiadm remove discovery-address`
- `iscsiadm remove static-config`
- `iscsiadm remove isns-server`

A user is no longer required to reboot a host when an attached iSCSI storage is removed from the host.

For further information, see the `iscsiadm(1M)` man page. See also the *System Administration Guide: Devices and File Systems*.

iSCSI MS/T Support

This system administration feature is new in the Solaris 10 6/06 release.

The iSCSI Multiple session target (MS/T) support feature enables a user to create more iSCSI session or paths to a target as needed. The additional iSCSI paths provide higher bandwidth aggregation and availability in specific configurations. The iSCSI MS/T support feature should be used in combination with MPxIO or other multipathing software.

The new `iscsiadm` commands are:

- `iscsiadm modify initiator-node -c number of sessions`
- `iscsiadm modify target-param -c number of sessions`

The iSCSI MS/T support feature enables higher bandwidth aggregation and availability to the customers with iSCSI arrays that support login redirection.

For more information, see:

- [iscsiadm\(1M\)](#) man page
- *System Administration Guide: Devices and File Systems*
- Using iSCSI Multipathing in the Solaris 10 Operating System <http://www.sun.com/blueprints/1205/819-3730.pdf>
- iSCSI RFC 3720 <http://www.ietf.org/rfc/rfc3720.txt?number=3720>

The `logadm` Utility

This system administration feature is new in the Solaris 10 6/06 release.

The `logadm` utility provides the `-l` option to rotate the logfile timestamps with local time. The `-l` option enables `logadm` to use the local time while naming files. This option does not change how timestamps are stored in log files.

For further information, see the [logadm\(1M\)](#) man page.

The `volfs` Utility

This enhancement to the `volfs` utility is new in the Solaris 10 6/06 release.

The volume management daemon, `vol`, is now managed by the Service Management Facility (SMF). This means you can use the `svcadm disable` command to disable the following new `volfs` service, if appropriate:

```
# svcadm disable volfs
```

You can identify the status of the `volfs` service by using this command:

```
$ svcs volfs
STATE          STIME      FMRI
online         Sep_29    svc:/system/filesystem/volfs:default
```

For more information, see the [smf\(5\)](#), [volfs\(7FS\)](#), and [vold\(1M\)](#) man pages.

For more information about managing the `volfs` service, see “What’s New in Removable Media?” in *System Administration Guide: Devices and File Systems*.

Basic Registration 1.1

This system administration feature is new in the Solaris 10 6/06 release.

Basic Registration 1.1 enables you to create a registration profile and ID to automate your Solaris software registrations. The software registration user interface and the procedure for registering Solaris software has been changed in Basic Registration 1.1.

For more information about the software registration user interface changes and step-by-step instructions about how to register your Solaris software, see the Sun Connection Information Hub at <http://www.sun.com/bigadmin/hubs/connection/>.

Sun Update Connection

This system administration feature is new in the Solaris 10 6/06 release.

Sun™ Update Connection System Edition is now known as Sun Update Connection. Sun Update Connection product supports the default `pat chadd` behavior from a global zone on a system where one or more non-global zones are installed.

For more information about the Sun Update Connection suite of products, see <http://docs.sun.com/app/docs/coll/1320.2>. See also the *Sun Update Connection, System Edition 1.0 Release Notes*.

Networking Enhancements

The following networking features and enhancements have been added to the Solaris 10 6/06 release.

IPv6 for IPFilter

This networking feature is new in the Solaris 10 6/06 release.

IPFilter for the Solaris OS has been updated to include IPv6 support. Packet filtering rules which include IPv6 addresses can be applied by using the `ipf` commands. IPv6 extension headers can be used to enable filtering. The IPv6 option has also been added to the `ipfstat` IPFilter statistics.

IPFilter can now be deployed in IPv6 networks to enhance security.

For further information see the `ipf(1M)` and `ipfstat(1M)` man pages. See also the *System Administration Guide: IP Services*.

UDP and TCP Performance Enhancement

This networking enhancement is new in the Solaris 10 6/06 release.

Performance of both the Transmission Control Protocol (TCP) and the User Datagram Protocol (UDP) has been enhanced in this release. The enhancements result in lower latency and higher throughput for both transmit performance and receive performance. Network applications yield better performance due to system performance improvements. This feature particularly benefits applications that heavily transmit and receive UDP packets or use TCP loopback connections.

For more information, see the [ip\(7P\)](#), [tcp\(7P\)](#), and [udp\(7P\)](#) man pages. See also the *Solaris Tunable Parameters Reference Manual*.

IP_NEXTHOP Socket Option

This networking enhancement is new in the Solaris 10 6/06 release.

The IP_NEXTHOP is a new IP-level socket option that specifies the address of the next hop for the traffic that originates from the socket. An application that has the IP_NEXTHOP option set bypasses the routing table lookups on the destination and sends packets directly to the specified `onlink nexthop`.

Note – The thread that sets the IP_NEXTHOP option must have the PRIV_SYS_NET_CONFIG privilege.

TCP_INIT_CWND Socket Option

This networking enhancement is new in the Solaris 10 6/06 release.

The new TCP socket option, TCP_INIT_CWND, enables an application to override the settings in the initial TCP congestion window, as described in RFC 3390, “Increasing TCP’s Initial Window.” By default, TCP sets the initial congestion window at connection set up time and after an idle period. (An idle period is when no traffic occurs between the two ends of the TCP connection. An application can use the TCP_INIT_CWND socket option to set the initial congestion window to a specified number of TCP segments. So the value of this new socket option is used both at the connection start time and after an idle period to set the initial congestion window. The process must have the PRIV_SYS_NET_CONFIG privilege if a number greater than that calculated by RFC 3390 must be specified.

For further information, see the [tcp\(7P\)](#) man page.

Security Enhancements

The following security features and enhancements have been added to the Solaris 10 6/06 release.

Enhancements for `pktool` Object Migration and Interoperability

These security enhancements are new in the Solaris 10 6/06 release.

The `pktool` command enables users to manage PKCS#11 objects. New subcommands have been added to move, display, and delete PKCS#11 objects and to show available PKCS#11 tokens. The new `pktool` subcommands help migrate cryptographic objects to or from the default Sun Software PKCS#11 Softtoken or other PKCS#11-compliant tokens.

For further information, see the [pktool\(1\)](#) man page.

SSL Proxy Module

This security enhancement is new in the Solaris 10 6/06 release.

A kernel-level SSL proxy server has been added in this release. The proxy simplifies and accelerates the SSL/TLS protocol implementation by pushing handshake and records processing to the kernel. The proxy supports the most commonly used cipher suites. You can configure applications, such as web servers, to offload the handling of the SSL operations with those cipher suites to the proxy, and seamlessly fall back to their existing user-level SSL library for the others.

For more information, see the [ksslcfg\(1M\)](#) man page.

AES Counter Mode

This security enhancement is new in the Solaris 10 6/06 release.

Advanced Encryption Standard (AES) is a block cipher recommended by the National Institute of Standard and Technology (NIST). When used in counter mode, a counter block is encrypted and the result is XOR with a block of plain text to produce cipher text. The counter mode is useful with block devices because the encryption or decryption of a block does not depend upon the prior encryption or decryption of any other block. The counter mode has been approved by NIST. This feature is only available to kernel consumers.

For more information, see the [libpkcs11\(3LIB\)](#) man page.

PKCS #11 v2.20 Support in the Solaris Cryptographic Framework

This security enhancement is new in the Solaris 10 6/06 release.

This feature adds RSA PKCS #11 v2.20 support to the Solaris Cryptographic Framework, including the stronger SHA2 algorithms.

For a list of mechanisms that v2.20 provides, see the [pkcs11_softtoken\(5\)](#) man page. For a list of mechanisms that are available to users, see the [digestp\(1\)](#) and [mac\(1\)](#) man pages.

Kerberos Cred Auto-Renew

This security enhancement is new in the Solaris 10 6/06 release.

In the Solaris 10 6/06 release, the Kerberos Cred Auto-Renew feature can automatically renew an user's credentials rather than sending a warning. The user is also not required to manually renew the credentials using the `kinit -R` command.

For more information, see [kttk_warnd\(1M\)](#) and [warn.conf\(4\)](#) man pages.

Device Management Enhancements

The following device management features and enhancements have been added to the Solaris 10 6/06 release.

iSNS Client Support for iSCSI

This device management feature is new in the Solaris 10 6/06 release.

The Internet Storage Name Service (iSNS) client feature adds a new discovery option to the Solaris OS iSCSI software initiator. This option enables user to use iSNS to handle Internet Protocol (IP)-SAN device discovery. This feature is off by default and is not a platform-specific feature. The iSNS client feature introduces several `iscsiadm` command additions, modifications, and driver changes to handle iSNS discovery.

IT customers who use iSCSI to build block-based IP-SAN need a scalable way to manage device discovery and configuration for their SANs as they grow. The iSNS client feature supports a scalable method for device discovery in a large IP-SAN configuration that uses a minimal configuration.

For further information about the new and modified command-line options, see the [iscsiadm\(1M\)](#) man page. See also the *System Administration Guide: Devices and File Systems*.

cdrecord, readCD, and cdda2wav Available

This device management feature is new in the Solaris 10 6/06 release.

Previously, `cdrecord` was available on a companion CD. In this release, `cdrecord` is now available in the Solaris OS. `cdrecord` is a powerful tool for burning CDs. `cdrecord` supports more burners than `cdrw`. `cdrecord` works better with USB and 1394 external burners. `cdrecord` is, however, limited to DVD images that are less than 2 Gbytes.

For further information, see the `cdrecord`, `readCD`, and `cdda2wav` man pages.

x86: PCI Express Support on x86 Systems

This device management feature is new in the Solaris 10 6/06 release.

This Solaris release provides support for the PCI Express (PCIe) interconnect. PCIe is designed to connect peripheral devices to desktop, enterprise, mobile, communication, and embedded applications.

The PCIe interconnect is an industry-standard, high-performance, serial I/O bus. For details on PCIe technology, go to the following site:

<http://www.pcisig.com>

The PCIe software provides the following features in this Solaris release:

- Support for extended PCIe configuration space
- Support for PCIe baseline error handling and MSI interrupts
- Modified IEEE-1275 properties for PCIe devices
- PCIe hotplug support (both native and ACPI-based) by enhancing the `cfgadm_pci` component of the `cfgadm` command

- ATTN button usage-based PCIe peripheral autoconfiguration

The following `cfgadm` example output displays the hotpluggable PCIe devices on an x86 system. Note that the display below might differ from platform to platform. Check your hardware platform guide for the correct `cfgadm` syntax.

```
# cfgadm pci
Ap_Id  Type          Receptacle  Occupant    Condition
pcie1  unknown      empty       unconfigured unknown
pcie2  unknown      empty       unconfigured unknown
pcie3  unknown      empty       unconfigured unknown
pcie4  etherne/hp   connected   configured  ok
pcie5  pci-pci/hp   connected   configured  ok
pcie6  unknown      disconnected unconfigured unknown
```

The administrative model for hotplugging PCIe peripherals is the same as for PCI peripherals, which use the `cfgadm` command.

For more information, see the `cfgadm_pci(1M)` man page and the *System Administration Guide: Devices and File Systems*. Check your hardware platform guide to ensure that PCIe and PCIe hotplug support is provided on your system. In addition, carefully review the instructions for physically inserting or removing adapters on your system. Also review the semantics of device auto-configuration, if applicable.

Solaris support for LSISAS1064 RAID Operations

This device management feature is new in the Solaris 10 6/06 release.

The `raidctl` utility creates, deletes, and displays Redundant Array of Inexpensive Disks (RAID) volumes of the LSI1030 and LSI1064 controllers. This utility also updates firmware, `fcode`, and BIOS for LSI1030 and LSI1064 controllers. The `raidctl` utility requires privileges controlled by the underlying file system permissions. Only privileged users can change the RAID system configuration. If a non-privileged user attempts to create or delete a RAID volume, the command fails with `EPERM`.

The LSI1030 SCSI Host Bus Adapter (HBA) supports a single RAID volume. This volume is a two-disk mirror, which is also known as Integrated Mirror (IM). IM volume is a RAID 1 type volume. An IM volume can be created on an LSI1030 controller only if the secondary disk has a capacity greater than or equal to the primary disk and there are no file systems mounted on the secondary disk.

The LSI1064 HBA can enable up to two RAID volumes, IM and Integrated Stripe (IS). The IS volume is a RAID 0 type volume. To create a volume on the LSI1064 controller, member disks of the volume should not have mounted file systems because volume initialization destroys data on member disks.

For more information, see `raidctl(1M)` man page.

Desktop Enhancements

The following desktop features and enhancements have been added to the Solaris 10 6/06 release.

32-bit: Palm Synchronization with USB Port

This desktop feature is new in the Solaris 10 6/06 release.

The palm synchronization feature enables Palm device synchronization through USB ports on the Solaris OS. This feature provides support to synchronize mobile devices such as Palm tops with desktops.

For further information, see the `gpilotd-control-applet(1)` man page.

32-bit: The `gnome-pilot` Utility

This desktop feature is new in the Solaris 10 6/06 release.

The `gnome-pilot` utility enables Palm users to synchronize their calendars, contacts and task lists between Evolution and their devices on the Solaris OS. This feature also enables the Palm user to back up and restore their devices on the Solaris OS. The `gnome-pilot` feature supports synchronization with the PalmOS® 4.x and the PalmOS 5.x.

x86: Xorg X Server Version 6.9

This X11 windowing feature enhancement is new in the Solaris 10 6/06 release.

The Xorg X Server for x86 and x64 platforms has been upgraded from version 6.8.2 to version 6.9 by the X.Org Foundation and the open source community. The new version adds support for more graphics devices including new models from ATI, XGI, VIA, and Intel. The Xorg X Server version 6.9 also adds keyboard and mouse handling improvements, performance enhancements, and bug fixes.

For further information, see the `Xorg(1)` man page.

Developer Tools Enhancements

The following developer tool enhancement has been added to the Solaris 10 6/06 release. To view developer tools that are new in the previous Solaris 10 1/06 release, see [“Developer Enhancements” on page 152](#).

`mediaLib 2.4`

`mediaLib 2.4` includes the following new features:

- Libraries have been tuned for the UltraSPARC-processor.
- A multithreaded (MT) library for better performance on multiprocessor x86 systems

For more information, see the [libmli\(3LIB\)](#) and [libmli_mt\(3LIB\)](#) man pages. See also <http://www.sun.com/processors/vis/mlib.html>.

New and Updated Drivers

The following drivers have been added or enhanced in the Solaris 10 6/06 release.

x86: SATA HBA Framework Support

This driver is new in the Solaris 10 6/06 release.

The Serial Advanced Technology Attachment (SATA) HBA framework project delivers a generic SATA framework for the Marvell 88SX60xx, Marvell 88SX50xx, and the Silicon Image 3124 controllers. The new SATA HBA drivers and framework provide native support for accessing SATA controllers and disks. These drivers provide features, such as hot plugging and queuing of multiple commands that are unique to SATA. These SATA drivers interface with the `sd` (target disk) driver and support the `sd` functionality.

For more information, see the following man pages:

- [sata\(7D\)](#)
- [cfgadm_sata\(1M\)](#)
- [si3124\(7D\)](#)
- [marvell88sx\(7D\)](#)
- [cfgadm\(1M\)](#)

USB-to-Serial Driver for Prolific Adapters

This driver is new in the Solaris 10 6/06 release.

This USB-to-serial driver supports Prolific pl2303 chipset-based adapters. With this new driver, customers can choose between Edgeport adapters and Prolific adapters.

For more information, see [usbsprl\(7D\)](#) man page.

Driver for IEEE 1394-Based (IIDC) Digital Cameras

This driver is new in the Solaris 10 6/06 release.

This driver enables support of IEEE 1394-based (IIDC) digital cameras. By supporting a software interface for camera control and image capture, the driver also enables development of applications that attach to these devices. This driver supports devices that implement the 1.04 version of the 1394 Trade Association 1394-based Digital Camera Specification. Support also covers devices that are backwards compatible.

For more information, refer to the [dcam1394\(7D\)](#) man page.

Driver Support for the New STK Tape Drive 10000 "Titanium"

This driver support is new in the Solaris 10 6/06 release.

This release enables the `st` (SCSI tape) driver to support the new StorageTek tape drive T10000A Titanium.

For further information, see the [st\(7D\)](#) man page.

USB-to-Serial Driver for Keyspan Adapters

This driver is new in the Solaris 10 6/06 release.

A new driver is provided in the release for Keyspan USB-to-serial adapters. This driver supports the USA-19HS model. With this new driver, customers have a choice between Edgeport adapters and Keyspan adapters.

For further information, see the [usbksp\(7D\)](#) man page.

Deimos Cryptographic Accelerator

This driver support is new in the Solaris 10 6/06 release.

The DCA driver supports Sun's SCA1000 and SCA500 cryptographic accelerator cards. Additionally, the driver will support Broadcom's 5820, 5821, and 5822 cards.

All of these cards support the following operations:

- RSA
- DSA
- 3DES
- DES
- RNG

The driver acts as cryptographic service provider to the Solaris Cryptographic Framework. The driver can be used by anyone who uses this framework.

x86: Driver Support for AMD64 Platform

This driver support is new in the Solaris 10 6/06 release.

In this release, the `glm` driver has been ported to the x64 platform. This driver enables you to use the x4422a card on the AMD64 platform.

For further information, see the [glm\(7D\)](#) man page.

rge Driver

This driver support is new in the Solaris 10 6/06 release. The `rge` driver provides support for Realtek RTL8169S/8110S Gigabit Ethernet devices.

For further information, refer to the [rge\(7D\)](#) man page.

Chelsio NIC Driver Support

In the Solaris 10 6/06 release, the `chxge` driver supports the Chelsio 10G Ethernet controller card. This support is provided on 32-bit and 64-bit architectures for both x86 and SPARC platforms. This driver supports the DLPI interface and checksum offload.

For detailed information, see the [chxge\(7D\)](#) man page.

HBA Drivers

These driver enhancements have been added to the Solaris 10 6/06 release.

The following HBA drivers have been added to the Solaris OS to support the full family of QLogic and Emulex HBA products. These products include both Sun and non-Sun HBAs:

- Solaris QLC common driver for Sun and QLogic branded HBAs
- Solaris EMLXS common driver for Sun and Emulex HBAs

These HBA drivers provide a selection of choices with a single fibre-channel implementation. The supported HBAs are Solaris Ready-certified and enable PCI-X and PCIe 4-Gbyte HBA support.

For further information, refer to the following links:

- http://www.sun.com/storage/san/infrastructure/fc_hba/
- <http://www.sun.com/solarisready>
- <http://qlogic.com>

Language Support Enhancements

The following language support features and enhancements have been added to the Solaris 10 6/06 release.

New UTF-8 Locales

This language support enhancement is new in the Solaris 10 6/06 release.

Over 50 new UTF-8 locales have been introduced in this release. As a result, Unicode support is now provided for all European, Middle Eastern, and Asian (EMEA) locales which lack a UTF-8 charset variant. Furthermore, locales for Cyprus, Luxembourg, and Malta are now available for the first time. Therefore, locales for all 25 European Union (EU) member states are now supported.

Additional Software

The following software have been added to the Solaris 10 6/06 release.

Watchdog Timeout

This additional software is new in the Solaris 10 6/06 release.

The watchdog mechanism detects a system hang. The watchdog timeout is a timer that is continually reset by a user application while the operating system and user application are running. While the watchdog timer operates in application mode, an additional alarm function, Alarm 3, is available to generate alarms for critical problems in the user application.

32-bit: RealPlayer for the Solaris OS

This additional software is new in the Solaris 10 6/06 release.

RealPlayer enables users to access and manage digital media. RealPlayer supports the following digital media formats:

- RealAudio
- RealVideo
- MP3
- Ogg Vorbis and Theora
- H263
- AAC

The RealPlayer feature enhances the multimedia experience for Sun Java Desktop System (Java DS) users.

pilot-link Software

This Open Source software is new in the Solaris 10 6/06 release.

Pilot-link is a suite of tools that enables you to connect your Palm or PalmOS compatible handheld with Unix, Linux, and any other POSIX-compatible machine. pilot-link works with almost all PalmOS handhelds. To sync Palm devices with Solaris by USB port, pilot-link makes use of `libusb`. For further information, see the [libusb\(3LIB\)](#).

This release of pilot-link is based on pilot-link v0.12.0-pre4.

For further information, see `pilot-xfer(1)` in the `/usr/sfw/man` directory.

PostgreSQL for the Solaris OS

This additional software is new in the Solaris 10 6/06 release.

PostgreSQL is a relational database system provided in the open-source community. More than 15 years of active development and a proven architecture has earned PostgreSQL a reputation for reliability, data integrity, and correctness.

For more information, see <http://www.postgresql.org>.

What's New in the Solaris 10 1/06 Release

This chapter summarizes all the new features in the Solaris 10 1/06 release.

Installation Enhancements

The following installation enhancements are new in the Solaris 10 1/06 release. To view installation features that are new in the previous Solaris 10 3/05 release, see [“Installation Enhancements” on page 186](#).

Upgrade Support Changes for Solaris Releases

Starting with the Solaris 10 1/06 release, you can upgrade the Solaris OS from the Solaris 8, 9, or 10 release. Upgrades from the Solaris 7 release are not supported. For further information, see the *Solaris 10 Installation Guide: Custom JumpStart and Advanced Installations*.

Sun Update Connection, System Edition 1.0

This feature is new in the Solaris 10 1/06 release.

The Sun Update Connection services enable you to keep your Solaris system up-to-date by giving you access to all the latest fixes and features. You can choose to perform local single-system update management or remote multiple-system update management. The Sun Update Connection services include these components:

- **Sun Update Manager.** The Sun™ Update Manager graphical user interface and the `smpatch` command-line interface enable you to manage updates locally on your Solaris system. The Sun Update Connection, System Edition software has the same functionality as the Sun Patch Manager tools, with the addition of some new features and enhancements.
- **Sun Update Connection.** This web application is hosted at Sun and enables you to manage updates remotely on one or more of your Solaris systems.
- **Sun Update Connection proxy.** This local caching proxy serves updates from Sun to the proxy's client systems inside your enterprise security domain.

- **SunSolve Patch & Updates Portal.** This portal gives you manual access to patches for download and to patch README files.

To receive the Sun Update Connection services, you must register your system. The registration process is initiated the first time that you boot your system or the first time that you access the Update Manager.

A subset of the Sun Update Connection services is available to Solaris users who do not have a service plan. These services include the use of Sun Update Manager to manage all security fixes and device driver updates that apply to your Solaris system.

To use the full suite of Sun Update Connection services, you must have one of the service plans described at <http://www.sun.com/service/solaris10/>. The full suite of services includes access to all patches, the Sun Update Connection web application, and the Sun Update Connection proxy.

For information about the Sun Update Connection, System Edition, see these books in the <http://docs.sun.com/app/docs/coll/1320.2/> on docs.sun.comSM.

- *Sun Update Manager 1.0 Administration Guide*
- *Sun Update Connection 1.0 Administration Guide*
- *Sun Update Connection, System Edition 1.0 Release Notes*

For information about registering your system after you install the Solaris software, see the Sun Connection Information Hub <http://www.sun.com/bigadmin/hubs/connection/>.

x86: GRUB Based Booting

Starting with the Solaris 10 1/06 release, the open source GNU GRand Unified Bootloader (GRUB) has been adopted in the Solaris OS for x86 based systems. GRUB is responsible for loading a boot archive into the system's memory. A boot archive is a collection of critical files that is needed during system startup before the root (/) file system is mounted. The boot archive is used to boot the Solaris OS.

The most notable change is the replacement of the Solaris Device Configuration Assistant with the GRUB menu. The GRUB menu facilitates booting the different operating systems that are installed on your system. The GRUB menu is displayed when you boot an x86 based system. From the GRUB menu, you can select an OS instance to install by using the arrow keys. If you do not make a selection, the default OS instance is booted.

The GRUB based boot feature provides the following improvements:

- Faster boot times
- Installation from USB CD or DVD drives
- Ability to boot from USB storage devices
- Simplified DHCP setup for PXE boot (no vendor-specific options)
- Elimination of all realmode drivers

- Ability to use Solaris Live Upgrade and the GRUB menu to quickly activate and fall back to boot environments

For more information about GRUB refer to the following sections.

Task	For More Information
Overview information about GRUB based booting and installing	<p>“GRUB Based Booting (Overview)” in <i>Solaris 10 Installation Guide: Solaris Live Upgrade and Upgrade Planning</i></p> <p><i>System Administration Guide: Basic Administration</i></p> <p><i>System Administration Guide: Devices and File Systems</i></p>
How to boot and install over the network with the GRUB menu	<p>“Booting and Installing the System From the Network With a DVD Image” in <i>Solaris 10 Installation Guide: Network-Based Installations</i></p>
How to boot and install with the GRUB menu and the Custom JumpStart installation method	<p>“Performing a Custom JumpStart Installation” in <i>Solaris 10 Installation Guide: Custom JumpStart and Advanced Installations</i></p>
How to use the GRUB menu and Solaris Live Upgrade to activate and fall back to boot environments	<p>“Activating a Boot Environment” in <i>Solaris 10 Installation Guide: Solaris Live Upgrade and Upgrade Planning</i></p> <p>Chapter 10, “Failure Recovery: Falling Back to the Original Boot Environment (Tasks),” in <i>Solaris 10 Installation Guide: Solaris Live Upgrade and Upgrade Planning</i></p>
How to perform system administration tasks with the GRUB menu	<p><i>System Administration Guide: Basic Administration</i></p> <p><i>System Administration Guide: Devices and File Systems</i></p> <p><code>bootadm(1M)</code></p> <p><code>installgrub(1M)</code></p>

Note – GNU is a recursive acronym for “GNU’s Not UNIX.” For more information, go to <http://www.gnu.org>.

Upgrading the Solaris OS When Non-Global Zones Are Installed

This feature is available in the Solaris 10 1/06 release.

The Solaris Zones feature provides the ability to configure non-global zones in a single instance of Solaris, the global zone. A non-global zone is an application execution environment in which processes are isolated from all other zones. If you are running a system with non-global zones installed, you can use standard Solaris upgrade programs to upgrade to the Solaris 1/06 release. You can use either the Solaris interactive installation program or custom JumpStart to upgrade. The process of upgrading with non-global zones installed has some limitations.

- A limited number of custom JumpStart keywords is supported.
- You must use the Solaris Operating System DVD or a DVD-created network installation image. You cannot use the Solaris Software CD media or a CD network installation image to upgrade a system.
- On a system with non-global zones installed, do not use Solaris Live Upgrade to upgrade your system. While you can create a boot environment with the `lu create` command, the `lu upgrade` command cannot upgrade a boot environment that has non-global zones installed. In that case, the upgrade fails and an error message is displayed.
- For a list of supported custom JumpStart keywords, see the *Solaris 10 Installation Guide: Custom JumpStart and Advanced Installations*.
- For details about using the Solaris interactive installation program, see the *Solaris 10 Installation Guide: Solaris Live Upgrade and Upgrade Planning*.

Networking Enhancements

The following networking enhancements are new in the Solaris 10 1/06 release. To view networking features that are new in the previous Solaris 10 3/05 release, see “[Networking Enhancements](#)” on page 222.

Source-Filtered Multicasting

This enhancement is new in the Solaris Express 6/05 release and in the Solaris 10 1/06 release.

Enhancements have been made to the Multicast Listener Discovery (MLD) protocol, for IPv6, and the Internet Group Management Protocol (IGMP), for IPv4. The Solaris implementation of these protocols has been enhanced to support MLDv2 and IGMPv3. These extensions provide support for source address filtering on multicast traffic. Also included is support for IETF-specified socket extensions. This support enables applications to take advantage of source-filtered multicasting.

For further information, see the *Programming Interfaces Guide* and the *System Administration Guide: IP Services*.

Enhancements for bge and xge Network Interfaces

This enhancement is new in the Solaris Express 6/05 release and in the Solaris 10 1/06 release.

The Solaris OS now contains features that extend the configuration possibilities for the bge and xge network interfaces. System administrators can now group these interfaces into LACP-capable link aggregations. These aggregations can support large-scale high availability or database implementations. In addition, you can configure xge and bge interfaces into virtual local area networks (VLANs) to extend network capability.

The new `dladm` command has been added for configuring and administering bge and xge interfaces. For further information, refer to the `dladm(1M)` man page.

Enhancements to Java Desktop System Release 3

The Solaris 10 1/06 release includes the Sun Java Desktop System Release 3 (Java DS). This section describes new and enhanced features in the Java DS.

New Features in Java DS

This release of the Java Desktop System on the Solaris 10 OS combines innovations from the open-source software community and from Sun Microsystems to create a comprehensive, integrated, and interoperable desktop environment. The following section describes improvements and enhancements in this release of the Java Desktop System.

Mobility Enhancements

The Java Desktop System has improved profile management for roaming web browser profiles.

Interoperability Enhancements

The Java Desktop System has the following interoperability enhancements.

- An Exchange connector enables the Email and Calendar applications to interoperate with a Microsoft Exchange network.
- Browser features improve compatibility with nonstandard browsers, such as Internet Explorer.
- Samba 3.0 provides better interoperability with Windows file sharing.

New Accessibility Applications

The Java Desktop System has the following accessibility enhancements.

- Screen reader and screen magnifier
- On-Screen Keyboard

See the *Java Desktop System Release 3 Accessibility Guide* and the *Java Desktop System Release 3 Accessibility Release Notes* for more details.

Internationalization Enhancements

The Java Desktop System has the following internationalization enhancements.

- Simplified Chinese SunPinYin input method
- New File System Examiner to enable migration from legacy environments to UTF-8 character encoding

General Improvements

The Java Desktop System has the following general enhancements.

- All of the major components of the Java Desktop System, GNOME, Email and Calendar, Mozilla™ browser, and Gaim have been updated.
- StarOffice™ 7 software includes Product Patch 4.
- The framework for Smartcard integration through industry-standard PC/SC API is included.

Plug-ins for Mozilla 1.7

This section contains information about plug-ins for Mozilla 1.7.

Java Plug-in

Java Plug-in is enabled in Mozilla 1.7 by default. If Java Plug-in does not work, then create a symbolic link in the `/usr/sfw/lib/mozilla/plugins` directory, pointing to the appropriate file shown in the following table.

Platform	Path
SPARC	<code>\$JAVA_PATH/plugin/sparc/ns7/libjavaplugin_oji.so</code>
x86	<code>\$JAVA_PATH/plugin/i386/ns7/libjavaplugin_oji.so</code>

In the Java Desktop System Release 3, `$JAVA_PATH` is `/usr/java/jre`.

Note – Do not copy the `libjavaplugin_oji.so` file instead of creating a symbolic link. A copy of the file causes the Java software to crash.

Install all Java plugins in the `/usr/sfw/lib/mozilla/plugins` directory.

Other Plugins

The Adobe Acrobat Reader third-party plugin is also available for Mozilla 1.7.

For more installation information about Mozilla plugins, see the following location:

<http://plugindoc.mozdev.org/linux.html>.

Localization Support

This release of the Java Desktop System includes localization enhancements.

Supported Languages

The supported languages for this release of the Java Desktop System are as follows:

- English
- French
- German
- Italian
- Japanese
- Korean
- Spanish
- Swedish
- Simplified Chinese
- Traditional Chinese
- Brazilian Portuguese – User interface support only

Note – This release of the Java Desktop System includes some localizations that are taken from the GNOME community. Sun Microsystems does not take any responsibility for the completeness or accuracy of these localizations.

Migration To Unicode Multilingual Computing

The Java Desktop System is a fully Unicode-enabled, multilingual system that supports languages with Unicode UTF-8 encoding. The Java Desktop System also provides codeset conversion to support legacy (non-UTF-8) encodings. See [Solaris 10 Release Notes](#) for recommended locales and descriptions of issues you might encounter when you migrate to Unicode multilingual computing.

Security Enhancements

The following security features are new in the Solaris 10 1/06 release. To view security features that are new in the previous Solaris 10 3/05 release, see [“Security Enhancements” on page 195](#).

Set SMTP to Use Transport Layer Security

This enhancement is new in the Solaris Express 8/05 release and in the Solaris 10 1/06 release.

The Simple Mail Transfer Protocol (SMTP) can use Transport Layer Security (TLS) in version 8.13 of sendmail. When enabled, this service to SMTP servers and clients provides private, authenticated communications over the Internet, as well as protection from eavesdroppers and attackers.

For more information, see the [System Administration Guide: Network Services](#).

Metaslot in the Cryptographic Framework

This feature is new in the Solaris 10 1/06 release and in the Solaris Express 2/05 release. This feature is of interest to both system administrators and software developers.

The metaslot is a component of the Solaris cryptographic framework library, `libpkcs11.so`. With metaslot software, an application that needs encryption can specify its cryptographic needs. With these specifications, the most suitable cryptographic mechanism that is available on the system will be supplied. The metaslot serves as a single virtual slot with the combined capabilities of all tokens and slots that have been installed in the framework. Effectively, the metaslot enables an application to connect transparently with any available cryptographic service through a single slot.

The metaslot is automatically enabled. The system administrator can explicitly disable the metaslot if preferred.

When an application requests a cryptographic service, the metaslot points to the most appropriate slot, which simplifies the process of selecting a slot. In some cases, a different slot might be required, in which case the application must perform a separate search explicitly.

Further information about the cryptographic framework is provided in the *Solaris Security for Developers Guide*. See also the *System Administration Guide: Security Services*.

IKE Enhancements

These enhancements are new in the Solaris 10 1/06 release and in the Solaris Express 2/05 release.

IKE is fully compliant with NAT-Traversal support as described in RFC 3947 and RFC 3948. IKE operations use the PKCS #11 library from the cryptographic framework, which improves performance. The cryptographic framework provides a `softtoken` keystore for applications that use the metaslot. When IKE uses the metaslot, you have the option of storing the keys on an attached board or in the `softtoken` keystore.

For further information about IKE, see the *System Administration Guide: IP Services*.

New Command for `embedded_su`

This enhancement is new in the Solaris 10 1/06 release.

This release includes a new command, `embedded_su`. This command offers “su-like” features to programs. This command enables graphical user interfaces to prompt for authentication data and execute operations as another user. This command provides functionality exactly equivalent to the `su` command, and so poses no new security risks. System administrators who have used `/etc/pam.conf` to customize the behavior of the `su` command can choose to add `/etc/pam.conf` entries to control the `embedded_su` command.

For an example, see the `embedded_su(1M)` man page.

System Performance Enhancements

The following system administration tools have been enhanced in the Solaris 10 1/06 release. To view performance enhancements that are new in the previous Solaris 10 3/05 release, see [“System Performance Enhancements” on page 211](#).

Large Pages Enhancement

This enhancement is new in the Solaris Express 6/05 release and in the Solaris 10 1/06 release.

This feature brings the benefits of large pages to a broader range of applications without requiring application tuning or system tuning. This enhancement applies large pages automatically for anonymous memory that is based on segment sizes.

Kernel Page Relocation

This enhancement is new in the Solaris Express 6/05 release and in the Solaris 10 1/06 release.

In these releases, system performance has been enhanced on the Sun Fire™ 15K, Sun Fire 20K, and Sun Fire 25K domains.

Memory Placement Optimization Hierarchical Lgroup Support

This feature is new in the Solaris 10 1/06 release and in the Solaris Express 6/05 release.

Hierarchical Lgroup Support (HLS) improves the Memory Placement Optimization (MPO) feature in the Solaris OS. HLS enables the Solaris OS to optimize performance for machines that have more than local memory and remote memory latencies. Machines with four or more Opteron CPUs might have local memory, remote memory, and further remote memory. On such machines, HLS enables the Solaris OS to distinguish between the degrees of remoteness. HLS enables the Solaris OS to allocate resources with the lowest possible latency for applications. The Solaris OS allocates local resources for a given application. And, if local resources are not available by default, the Solaris OS allocates the nearest remote resources.

The *Programming Interfaces Guide* explains the abstraction that is used by the Solaris OS to identify which resources are near each other for optimizing localization. The guide also described the API that can be used for the locality group (lgroup) abstraction. Man pages are also available for this API. See `liblgrp(3LIB)`.

System Administration Enhancements

The following system administration tools have been enhanced in the Solaris 10 1/06 release. To view system administration enhancements that are new in the previous Solaris 10 3/05 release, see [“System Administration Enhancements” on page 160](#).

Improvements to Volume Management (vol0)

This volume management feature is new in the Solaris 10 1/06 release and in the Solaris Express 11/05 release.

The `vol` command is now hot-plug aware. This improvement means that if you insert removable media, the media is automatically detected and mounted by `vol`. You do not need to restart `vol` manually to recognize and mount a file system from any removable media device.

If you are using a legacy or non-USB diskette device, then you might need to issue the `vol check` command before `vol` can recognize the media. If the media is detected, but for some reason, is unmounted, then you'll need to run the following command:

```
# volrmmount -i rmdisk0
```

Before you hot-remove a removable media device, eject the media first. For example:

```
# eject rmdisk0
```

For more information about using `vol`, see [Chapter 1, “Managing Removable Media \(Overview\)”](#), in *System Administration Guide: Devices and File Systems*.

SunVTS 6.1

SunVTS 6.1 software is available in the Solaris 10 1/06 release. SunVTS 6.0 software is available in the previous Solaris 10 3/05 release.

The SunVTS software is a comprehensive software diagnostic package. SunVTS tests and validates Sun x86 and SPARC hardware. The test suite verifies the configuration and proper functioning of controllers, devices, and platforms.

SunVTS x86 diagnostics are now supported in the AMD 64-bit environment for the SunVTS kernel (`vt sk`). All x86 diagnostics except the System Test (`sys test`) are ported to the 64-bit platform.

New x86 diagnostics for this release include the following:

- Level 1 Data Cache Test (`l1dcachetest`)
- Infiniband Host Channel Adapter Test (`ibhctest`)

New SPARC diagnostics for this release include the following:

- Level 3 Cache Test (`l3sramtest`)
- Netra Intelligent Platform Management Controller Test (`nipmctest`)
- Loopback test support for Qlogic 2202 Board Test (`qlctest`)
- 64-bit x86 support for Tape Drive Test (`tapetest`)

Refer to the SunVTS 6.1 documentation at <http://docs.sun.com/app/docs/coll/sunvts6.1-s10-1-06>. This documentation provides details about these new features and tests.

Support for SCSI Disks Larger Than 2 Terabytes

This enhancement is new in the Solaris Express 6/05 release and in the Solaris 10 1/06 release.

SCSI, Fibre Channel, and iSCSI disks that are larger than 2 terabytes are now supported on 64-bit platforms. The format utility can be used to label, configure, and partition these larger disks.

For more information, see the *System Administration Guide: Devices and File Systems*.

Fibre-Channel HBA Port Utility

This feature is new in the Solaris 10 1/06 release and in the Solaris Express 2/05 release.

`fcinfo` is a command-line interface that collects administrative information about fibre-channel HBA ports. This interface also collects data about any Fibre-Channel targets that might be connected to those ports in a Storage Area Network (SAN).

For further information, see the `fcinfo(1M)` man page.

Additional Banner Page Printing Option in Solaris Print Manager

This enhancement is new in the Solaris 10 1/06 release and in the Solaris Express 3/05 release.

Solaris Print Manager has been expanded to include an additional “Never Print Banner” option. This option ensures that banner pages are never printed for the specified print queue.

Previously, you only had two choices for printing banner pages in Solaris Print Manager:

- You could enable the “always print banner” option in Solaris Print Manager.
- You could select the banner on or off option when you submitted a print job. This option was on by default.

The current printing options in the Print Manager reflect the `lpadmin` command options for printing to local print queues.

For further information about the Solaris Print Manager, see the *System Administration Guide: Advanced Administration*.

x86: New `prtconf` Option to Display Product Names

This enhancement is new in the Solaris Express 7/05 release and in the Solaris 10 1/06 release.

A new `-b` option has been added to the `prtconf` command. This option displays the product name of a system. This option is similar to the `uname -i` command. However, the `prtconf -b` command is specifically designed to determine the marketing name of a product.

The `-b` option displays the following root properties from the firmware device tree:

- `name`
- `compatible`
- `banner-name`
- `model`

To display additional platform-specific output that might be available, add the `-v` option to the `prtconf -b` command.

For more information, see the `prtconf(1M)` man page and the *System Administration Guide: Advanced Administration*.

Developer Enhancements

The following developer tool enhancements have been added to the Solaris 10 1/06 release. To view developer tools that are new in the previous Solaris 10 3/05 release, see “[Developer Tool Enhancements](#)” on page 235.

mediaLib 2.3

This feature is new in the Solaris Express 10/05 release and in the Solaris 10 1/06 release.

mediaLib is a low-level application library for building portable, high-performance, multimedia applications. These applications are useful in the telco, bioinformatics, and database fields. mediaLib 2.3 introduces a multithreaded (MT) library for better performance on multiprocessor systems. mediaLib 2.3 also introduces 340 new functions.

For further information, see the `libmLib(3LIB)` and `libmLib_mt(3LIB)` man pages. See also <http://www.sun.com/processors/vis/mlib.html>.

Advanced DDI Interrupts

This enhancement is new in the Solaris Express 6/05 release and the Solaris 10 1/06 release.

The Solaris OS now provides a new DDI Interrupt framework for registering interrupts and deregistering interrupts. Support for Message Signalled Interrupts (MSIs) is also provided. New management interfaces enable you to manipulate priorities and capabilities, to interrupt masking, and to obtain pending information.

The framework includes the following new interfaces:

- `ddi_intr_add_handler`
- `ddi_intr_add_softint`
- `ddi_intr_alloc`
- `ddi_intr_block_disable`
- `ddi_intr_block_enable`
- `ddi_intr_clr_mask`
- `ddi_intr_disable`
- `ddi_intr_dup_handler`
- `ddi_intr_enable`
- `ddi_intr_free`
- `ddi_intr_get_cap`
- `ddi_intr_get_hilevel_pri`

- `ddi_intr_get_navail`
- `ddi_intr_get_nintrs`
- `ddi_intr_get_pending`
- `ddi_intr_get_pri`
- `ddi_intr_get_softint_pri`
- `ddi_intr_get_supported_types`
- `ddi_intr_remove_handler`
- `ddi_intr_remove_softint`
- `ddi_intr_set_cap`
- `ddi_intr_set_mask`
- `ddi_intr_set_pri`
- `ddi_intr_set_softint_pri`
- `ddi_intr_trigger_softint`

Note – To take advantage of the features of the new framework, developers need to use the new interfaces. Avoid using the following interfaces, which are retained for compatibility only:

- `ddi_add_intr`
 - `ddi_add_softintr`
 - `ddi_dev_nintrs`
 - `ddi_get_iblock_cookie`
 - `ddi_get_soft_iblock_cookie`
 - `ddi_iblock_cookie`
 - `ddi_idevice_cookie`
 - `ddi_intr_hilevel`
 - `ddi_remove_intr`
 - `ddi_remove_softintr`
 - `ddi_trigger_softintr`
-

For more information, see “Interrupt Handlers” in the *Writing Device Drivers* manual. See also the individual man pages for the new interfaces. All man pages for these interfaces are in the 9F man page section.

Desktop Enhancements

The following desktop enhancements have been added to the Solaris 10 1/06 release. To view desktop enhancements that are new in the previous Solaris 10 3/05 release, see [“Desktop Enhancements” on page 176](#).

SPARC: Adobe Acrobat Reader 7.0.1

This enhancement is new in the Solaris Express 12/05 release and in the Solaris 10 1/06 release.

In this release, Adobe® Acrobat Reader is updated from version 5.0 to version 7.0.1. Adobe Reader enables you to view, navigate, and print Portable Document Format (PDF) files. This enhancement is for the SPARC platform.

SPARC: Sun OpenGL 1.5 for Solaris

This feature is new in the Solaris 10 1/06 release.

Sun OpenGL® for Solaris has been revised to cover the 1.5 specification for devices which can support OpenGL 1.5 functionality. Sun OpenGL 1.5 for the Solaris OS is available with the Solaris 10 1/06 release.

For further information, see the man pages are under the following directories:

- /usr/openwin/man/man3gl
- /usr/openwin/man/man3glu
- /usr/openwin/man/man3glw
- /usr/openwin/man/man3glx

Xorg Configuration GUI

This enhancement is new in the Solaris 10 1/06 release.

`xorgcfg`, the Xorg Configuration GUI, allows for easier configuration of the Xorg X server on x86 systems.

Xorg Release 6.8.2

This enhancement is new in the Solaris 10 1/06 release and in the Solaris Express 2/05 release.

The Xorg X server implementation has been upgraded from the 6.8.0 version to version 6.8.2 in this release. This upgrade fixes several bugs in the modules for various graphics cards. The upgrade also adds support for new graphics card models.

X Client Support for XFree86 Extensions

These X11 windowing enhancements are new in the Solaris Express 6/05 release and in the Solaris 10 1/06 release.

These enhancements take advantage of X server extensions that were originally developed by the XFree86 Project. These enhancements are now incorporated into the Xorg X server.

The new programs available include the following:

- | | |
|-----------------------|---|
| <code>xgamma</code> | Alters a monitor's gamma correction through XF86VidMode extension |
| <code>xrandr</code> | Resizes or rotates the screen through the RandR extension |
| <code>xvidtune</code> | Provides video mode tuner for Xorg through XF86VidMode extension |
| <code>xvinfo</code> | Prints X Video extension adapter information |

Advanced users can use these applications to tune the settings of the Xorg server at runtime. This process provides more information about the capabilities of the current system hardware.

Note – These programs require support for the XFree86 extensions. As such, these programs do not currently work with the Xsun X server. These programs do not work with other X servers that lack this support.

For more information, see the man pages for each application. Add `/usr/X11/man` to your `$MANPATH`, if needed, to view the man pages.

Support for Virtual USB Keyboard and Mouse Devices

This enhancement is new in the Solaris Express 6/05 release and in the Solaris 10 1/06 release.

This feature enables the Solaris OS to support multiple keyboards and mouse devices simultaneously. The `virtualkm` feature also provides an auto-switch function that enables users to obtain separate input for keyboards and for mouse devices. All these enhancements are compatible with existing applications.

This feature is particularly useful for the following users:

- Users of systems with KVMs support
- Laptop users who connect an extra mouse
- Users with special keypads or other devices that are declared as keyboards or mouse devices

For more information, see [virtualkm\(7D\)](#).

Device Management Enhancements

The following device management enhancements have been added to the Solaris 10 1/06 release. To view device management tools that are new in the previous Solaris 10 3/05 release, see [“Device Management” on page 213](#).

Support for iSCSI Devices

This feature is new in the Solaris 10 1/06 release and in the Solaris Express 2/05 release.

Support for Internet Small Computer System Interface (iSCSI) is provided in this release. iSCSI is an Internet Protocol (IP)-based storage networking standard for linking data storage subsystems. By carrying SCSI commands over IP networks, the iSCSI protocol enables you to mount disk devices, from across the network, onto your local system. On your local system, you can use the devices like block devices.

The iSCSI protocol does the following:

- Runs across existing Ethernet networks
- Uses existing management tools for IP networks
- Can be used to connect to Fibre-Channel or iSCSI Storage Area Network (SAN) environments

You can use the `iscsiadm` command to set up and manage your iSCSI devices. For more information, see the *System Administration Guide: Devices and File Systems* and the `iscsiadm(1M)` man page.

New and Updated Drivers

The following driver enhancements have been added to the Solaris 10 1/06 release. To view driver features that are new in the previous Solaris 10 3/05 release, see “[New or Updated Drivers](#)” on page 213.

SPARC: Sun XVR-2500 Graphics Accelerator

This driver is new in the Solaris 10 1/06 release.

This release introduces the `kfb` driver that supports the XVR-2500 graphic board. For further information, see the `SUNWkfb_config(1M)` and `kfb(7D)` man pages.

SCSI HBA Driver for LSI MegaRAID 320-2x SCSI RAID Controller

This driver is new in the Solaris Express 8/05 release and in the Solaris 10 1/06 release.

This release introduces the `lsimega` driver. The LSI 320-2x card is supported by `lsimega` with back-end support from LSI. Sun Fire V20z/V40z servers with the optional 320-2x RAID card receive better support for RAID disk I/O and passthrough I/O for SCSI tape and CDROM.

For further information, see the `lsimega(7D)` man page.

USB CCID IFD Handler

This driver is new in the Solaris Express 8/05 release and in the Solaris 10 1/06 release.

This release delivers a userland device driver for USB CCID-compliant smart card readers. This feature uses the USB CCID IFD Handler that is part of the MUSCLE project. With this feature integrated into Solaris, customers can now use a wide variety of USB smart card readers on Solaris systems that have USB ports.

For further information, see the `usb_ccid_ifd(3SMARTCARD)` man page.

More documentation is available at the public MUSCLE card web site at

<http://www.musclicard.com>.

New nge Driver

In the Solaris 10 1/06 release, the nge driver supports the NVIDIA CK8-04 gigabit Ethernet interface. See the nge(7D) man page.

Documentation Enhancements

The following documentation improvements have been made in the Solaris 10 1/06 release. To view documentation changes that were made in the previous Solaris 10 3/05 release, see “[Documentation Changes](#)” on page 263.

Solaris 10 1/06 Documentation

Solaris 10 documentation contains the Solaris 10 base documentation plus new documentation and revisions that have been published since the Solaris 10 3/05 release. New and revised documentation for the Solaris 10 1/06 release, therefore, can be found in the Solaris 10 collections that are available at <http://docs.sun.com/app/docs/prod/solaris.10>.

Sun Fire Documentation Changes

In the Solaris 10 1/06 release, three Sun Fire midrange and high-end system user guides have been combined into a single book. The new book is *Sun Fire High-End and Midrange Systems Dynamic Reconfiguration User Guide* (819-1501-10).

This new book replaces the following three previous documents:

Sun Fire Midrange Systems Dynamic Reconfiguration User Guide (817-7167-10)

Sun Fire High-End Systems Dynamic Reconfiguration User Guide (817-7166-10)

System Management Services 1.4 Dynamic Reconfiguration User Guide (817-4459-10 for Sun Fire High-End systems)

The new book is available in the *Solaris 10 on Sun Hardware Collection* at <http://docs.sun.com/app/docs/coll/1140.2>.

What's New in the Solaris 10 3/05 Release

This chapter summarizes all the new features in the Solaris 10 3/05 release. This chapter also summarizes all features, which are sorted by the Software Express release that introduced these features.

Key Features in Solaris 10 Software

In the Solaris 10 OS, Sun Microsystems has developed a new architecture for building and deploying systems and services that are capable of Predictive Self-Healing. Self-healing technology enables Sun systems and services to maximize availability. See [“Predictive Self-Healing” on page 160](#). The Predictive Self-Healing feature includes significant changes to the booting and service administration processes. See [“Solaris Service Manager” on page 160](#). Also in the Solaris 10 OS, changes in the installation of the operating system provide a simplified and unified installation process. See [“Solaris Installation Changes Including Installation Unification” on page 186](#).

The Solaris 10 OS introduces Solaris Zones software partitioning technology. Solaris Zones is a component of the Solaris Containers environment. Zones are used to virtualize operating system services and provide an isolated and secure environment for running applications. See [“Solaris Zones Software Partitioning Technology” on page 178](#).

Other key features in the Solaris 10 software include the DTrace dynamic tracing facility, process rights management, and a new architecture for network stacks. DTrace is a comprehensive dynamic tracing facility that gives Solaris users, administrators, and developers a new level of observability into the kernel and user processes. For further information, see [“DTrace Dynamic Tracing Facility” on page 162](#). In the Solaris software, processes that previously required superuser capabilities now require process rights. Process rights management uses privileges to restrict processes to only those rights that are required to perform the task. See [“Process Rights Management” on page 195](#) for further information. Also of particular importance, the networking stack for TCP connections has been rearchitected to deliver extremely high performance while improving the scalability. For further information, see [“New Architecture for Network Stacks” on page 211](#).

Java 2 Platform, Standard Edition 5 is another key technology. See [“Enhancements to the Java 2 Platform, Standard Edition 5” on page 183](#) for descriptions of the enhancements that are available in this platform. Also of particular significance, the Solaris 10 software introduces support for the 64-bit computing capabilities of the AMD Opteron processor. For details, see [“64-bit Support on x86 Systems” on page 163](#). Finally, the Solaris 10 software introduces the Java Desktop System, Release 3. This desktop system combines open source software with Sun innovation. For further information, see [“Java Desktop System, Release 3” on page 176](#).

These key features and many other new features in the Solaris 10 software are summarized in this book. Most features in the Solaris 10 3/05 release were introduced through the Software Express program. To view a list of all Solaris 10 features by their original Solaris Express release date, see [“Features by Software Express Release Date” on page 267](#).

System Administration Enhancements

This section describes all system administration features in the Solaris 10 3/05 release that are new or have been enhanced since the Solaris 9 OS was originally distributed in May 2002. DTrace and Predictive Self-Healing are features of particular importance. To view system administration tools that are new in the Solaris 10 7/05 release, see [“System Administration Enhancements” on page 149](#).

Predictive Self-Healing

This feature is new in the Solaris Express 6/04 release. The Solaris Express 10/04 release and the Solaris 10 3/05 release provided important enhancements.

Sun Microsystems has developed a new architecture for building and deploying systems and services capable of Predictive Self-Healing. Self-healing technology enables Sun systems and services to maximize availability when software and hardware faults occur. In addition, the self-healing technology facilitates a simpler and more effective end-to-end experience for system administrators and service providers, thereby reducing costs. The first major set of new features to result from this initiative is available in the Solaris 10 OS. The Solaris 10 software includes components that facilitate self-healing for CPU, memory, I/O bus nexus components, and system services.

Specific details about the components of this new architecture are covered in the following descriptions for the Solaris Service Manager and the Solaris Fault Manager.

Solaris Service Manager

Introduced in the Solaris Express 10/04 release and enhanced in the Solaris 10 3/05 release, the Solaris Service Manager provides an infrastructure that augments the traditional UNIX startup scripts, init run levels, and configuration files. This infrastructure provides the following functions:

- Automatically restarts failed services in dependency order, whether the services failed as the result of administrator error, a software bug, or an uncorrectable hardware error.

- Makes services objects that can be viewed, with the new `svcs` command, and managed, with `svcadm` and `svccfg` commands. You can also view the relationships between services and processes by using `svcs -p`, for both SMF services and legacy `init.d` scripts.
- Makes it easy to back up, restore, and undo changes to services by taking automatic snapshots of service configurations.
- Makes it easy to debug. You can ask questions about services and receive an explanation of why a service isn't running by using `svcs -x`. Also, this process is eased by individual and persistent log files for each service.
- Enhances the ability of administrators to securely delegate tasks to nonroot users, including the ability to modify properties and start, stop, or restart services on the system.
- Boots faster on large systems by starting services in parallel according to the dependencies of the services. The opposite process occurs during shutdown.
- Enables you to customize the boot console output either to be as quiet as possible, which is the default, or to be verbose by using `boot -m verbose`.
- Preserves compatibility with existing administrative practices wherever possible. For example, most customer and ISV-supplied `rc` scripts still work as usual.
- Enables you to configure your system services in one of two modes, both represented as `smf(5)` profiles. The “`generic_open.xml`” profile enables all the traditional Internet services that were previously enabled by default in the Solaris OS. The “`generic_limited_net.xml`” profile disables a large number of services that are frequently disabled during the process of hardening a system. However, this profile is not a replacement for the Solaris Security Toolkit (JASS) tool. See the individual profiles for details.

See Chapter 9, “Managing Services (Overview),” in the *System Administration Guide: Basic Administration* for more information about this infrastructure. An overview of the infrastructure can be found in the `smf(5)` man page.

Solaris Fault Manager

Predictive Self-Healing systems include a simplified administration model. Traditional error messages are replaced by telemetry events that are consumed by software components. The software components automatically diagnose the underlying fault or defect and initiate self-healing activities. Examples of self-healing activities include administrator messaging, isolation or deactivation of faulty components, and guided repair. A new software component is called Fault Manager, `fmd(1M)`. The Fault Manager manages the telemetry, log files, and components. The new `fmadm(1M)`, `fmdump(1M)`, and `fmstat(1M)` tools are also available in the Solaris 10 OS to interact with the Fault Manager and new log files.

When appropriate, the Fault Manager sends a message to the `syslogd(1M)` service to notify an administrator that a problem has been detected. The message directs administrators to a knowledge article on Sun's new message Web site, <http://www.sun.com/msg/>, which explains more about the problem impact and appropriate responses and repair actions.

The Solaris Express 6/04 release introduced self-healing components for automated diagnosis and recovery for UltraSPARC-III and UltraSPARC-IV CPU and memory systems. This release also provided enhanced resilience and telemetry for PCI-based I/O.

DTrace Dynamic Tracing Facility

DTrace is a feature of key importance in the Solaris 10 OS. DTrace is new in the Solaris Express 11/03 release and was enhanced in the Solaris Express 5/04 release and the Solaris Express 11/04 release.

DTrace is a comprehensive dynamic tracing facility that gives Solaris users, administrators, and developers a new level of observability into the kernel and user processes.

DTrace helps you understand your system by permitting you to dynamically instrument the operating system kernel and user processes to record data that you specify at locations of interest, called *probes*. Probes are like small programmable sensors that are scattered all over your Solaris system. More than 30,000 probes are included in this release. Each probe can be associated with custom programs that are written in the new D programming language. This association enables you to access system data by using ANSI C types and expressions and easily capture stack traces, record timestamps, build histograms, and more.

All of DTrace's instrumentation is entirely dynamic and available for use on your production system. When DTrace is off, you have no performance cost. The performance impact of any tracing is limited to only those probes and actions that you enable. Moreover, DTrace is safe: you can't damage the running system because DTrace has security, complete safety, and error checking at the core of its design. These features enable you to use DTrace with confidence on your running system whenever you need the facility to help investigate a problem.

The *Solaris Dynamic Tracing Guide* and a set of man pages, including `dttrace(1M)`, are provided to help you learn DTrace. The guide includes a complete feature reference and examples to help you get started. You can also join the DTrace forum on BigAdmin (see <http://www.sun.com/bigadmin/content/dttrace>) to learn more about DTrace, download example programs, and ask questions.

In the Solaris Express 5/04 release, user-process tracing that uses the `pid` provider was made available for x86 systems. This feature has been available on SPARC platforms since the introduction of DTrace in the Solaris Express 11/03 release.

The `pid` provider enables you to trace any instruction in any process either at the level of any function call's entry and return, or at any offset into any function. For complete details, see Chapter 27, "pid Provider," and Chapter 32, "User Process Tracing," in the *Solaris Dynamic Tracing Guide*.

In the Solaris Express 11/04 release, the `plockstat(1M)` utility enables you to observe user-level synchronization primitives, such as `lockstat(1M)` in the kernel. The DTrace `plockstat` provider is the underlying instrumentation methodology for `plockstat(1M)`. Dtrace `plockstat` can be used to augment the data recorded by the `plockstat` utility. See the `plockstat(1M)` man page for further information.

64-bit Support on x86 Systems

This feature is new in the Solaris Express 11/04 release.

The Solaris 10 release contains support for the 64-bit computing capabilities of AMD64 and EM64T platforms. This support includes the ability to run both 32-bit and 64-bit applications in the 64-bit computing environment.

For details about supported hardware configurations, see <http://www.sun.com/bigadmin/hcl>.

Note – The Solaris 10 software already supports these platforms for 32-bit computing.

For more information about Solaris for 64-bit support on x86 systems, see the following references in the Solaris 10 documentation:

- Appendix C, “Making a Device Driver 64-bit Ready” in *Writing Device Drivers*
- Chapter 9, “Shutting Down and Booting a System (Overview),” in *System Administration Guide: Basic Administration*
- *Solaris 64-bit Developer’s Guide*

Kernel Selection for x86 Systems

This feature is new in the Solaris Express 11/04 release.

The Solaris 10 OS has the ability to automatically detect whether your system is 64-bit capable and then boot the appropriate kernel.

Following a new installation of the Solaris 10 software, the boot program automatically loads the 64-bit kernel if it detects your system is 64-bit capable. Otherwise, the program loads the 32-bit kernel.

Following an upgrade installation of the Solaris 10 OS on a system that is configured to load the default 32-bit kernel, the system now automatically determines whether to load the 32-bit or 64-bit kernel. If the system was configured to load a non-default kernel, the system continues to load that non-default kernel. Procedures for customizing a system to load a specific kernel are outlined in Chapter 8, “Shutting Down and Booting a System (Overview)” in *System Administration Guide: Basic Administration*.

Further documentation about the kernel selection procedure is provided in the Solaris 10 documentation at <http://docs.sun.com>.

Sun Java Web Console

This feature is new in the Solaris Express 10/04 release.

The Sun Java Web Console provides a common location for users to access web-based management applications. Users reach the console by logging in through an HTTPS port, using one of several supported web browsers. The single entry point that is provided by the console eliminates the need to learn URLs for multiple applications. This entry point provides authentication and authorization for all applications that are registered with the console.

All console-based applications conform to the same user interface guidelines. This change enhances ease of use by enabling users to transfer their knowledge of one application to another.

The Java Web Console also provides auditing and logging services for all registered applications.

Note – Currently, no system management applications in the Solaris OS use the Java Web Console.

For more information about the Java Web Console, see the *System Administration Guide: Basic Administration*.

Multi-Owner Disk Set Support With Solaris Volume Manager for Sun Cluster

This feature is new in the Solaris Express 10/04 release and in the Solaris 9 9/04 release.

Solaris Volume Manager for Sun Cluster introduces volume management features that work specifically with Sun Cluster and with applications such as Oracle Real Application Clusters. Solaris Volume Manager for Sun Cluster enables you to create and manage storage that is grouped into multi-owner disk sets. A multi-owner disk set enables multiple nodes to share ownership of a disk set. Multi-owner disk sets provide scalability because an instance of the application runs on each node in the cluster. Because each instance of the application directly accesses the shared storage, multi-owner disk sets also enhance the performance of the application.

For more information, see the *Solaris Volume Manager Administration Guide*. For information about Sun Cluster, see *Sun Cluster Software Installation Guide for Solaris OS* and *Sun Cluster Data Service for Oracle Real Application Clusters Guide for Solaris OS*.

Expanded Disk Set Support in the Solaris Volume Manager

This feature is new in the Solaris Express 1/04 release and in the Solaris 9 9/04 release.

The Solaris Volume Manager has expanded the functionality of device IDs. In addition to tracking movement in local sets, the Solaris Volume Manager now also automatically tracks disk movement within named disk sets. This enhancement ensures that Solaris Volume Manager volumes remain intact even if disks are rearranged. This device ID support also enables two hosts to share access nonconcurrently to disk sets, even with dissimilar views of accessible raw storage.

The Solaris Volume Manager has also been enhanced with the `metaimport` command. This command uses the expanded device ID support to enable you to import disk sets, even disk sets that were created on different systems.

See the [`metaimport\(1M\)`](#) man page for further information.

Importing Remotely Replicated Disk Sets With Solaris Volume Manager

This feature is new in the Solaris Express 10/04 release.

Solaris Volume Manager has expanded the functionality of the `metaimport` command. In addition to importing regular disk sets from one system to another, you can now import replicated disk sets.

A replicated disk set is created by using remote replication software, such as Sun StorEdge™ Network Data Replicator (SNDR) software or Hitachi TrueCopy.

For more information, see the [`metaimport\(1M\)`](#) man page and the *Solaris Volume Manager Administration Guide*.

Device ID Changes

This description is new in the Solaris Express 10/04 release.

Device ID output is displayed in a new format in the Solaris Express 10/04 release. Previously, the device ID was displayed as a hexadecimal value. The new format displays the device ID as an ASCII string. Solaris Volume Manager might display the device ID output in the new or old format, depending on when the device id information was added to the state database replica.

For more information, see the *Solaris Volume Manager Administration Guide*.

Multiterabyte Volume Support in Solaris Volume Manager

Multiterabyte volume support is available only for systems that run a 64-bit kernel. This feature is new in the Software Express pilot program and in the Solaris 9 4/03 release. This feature is included in the Solaris 10 3/05 release.

Solaris Volume Manager has been enhanced with multiterabyte volume support. With this support, the Solaris Volume Manager can create, manage, and delete large (>1Tbyte) RAID-0 (stripe), RAID-1 (mirror), RAID-5, and soft-partition volumes. In addition, this large volume support enables Solaris Volume Manager to construct volumes on large or EFI-labeled logical unit numbers (LUNs).

Solaris Volume Manager large volume support is not available for systems running a 32-bit Solaris kernel.

See also “[Multiterabyte Disk Support With EFI Disk Labels](#)” on page 250.

For further information, see the *Solaris Volume Manager Administration Guide*.

Reconfiguration Coordination Manager Support in Solaris Volume Manager

This feature is new in the Software Express pilot program and in the Solaris 9 4/03 release. This feature is included in the Solaris 10 3/05 release.

Reconfiguration Coordination Manager (RCM) support adds to Solaris Volume Manager the ability to respond appropriately to dynamic reconfiguration (DR) requests. This addition ensures that removal of devices under Solaris Volume Manager control is blocked with an appropriate warning. This block remains in effect until the devices are no longer in use. This warning prevents system administrators from accidentally removing active volumes from a DR-configured system.

For further information, see the [Solaris Volume Manager Administration Guide](#).

Top-Down Volume Creation in Solaris Volume Manager

This feature is new in the Solaris Express 9/03 release and in the Solaris 9 4/04 release.

The top-down volume creation feature of Solaris Volume Manager provides a new command that enables system administrators to quickly and easily create quality of service-based Solaris Volume Manager configurations. Rather than manually partitioning disks, creating stripes, and assembling mirrors, the `metassist` command manages the details and provides functional logical volumes. The volumes are based on the criteria that are specified on the command line or in referenced configuration files.

For further information, see the `metassist(1M)` man page and the [Solaris Volume Manager Administration Guide](#).

Package and Patch Tool Enhancements

This feature is new in the Solaris Express 11/03 release under the title “Solaris Product Registry Enhancements.” This feature description was updated in the Solaris Express 8/04 and 10/04 releases.

The Solaris package and patch tools have been enhanced, providing improved performance and extended functionality.

As a part of these enhancements, the `pkgchk` command now provides a new option to assist you in mapping files to packages. To map files to packages, use the `pkgchk -P` option instead of `grep pattern /var/sadm/install/contents`. The `-P` option enables you to use a partial path. Use this option with the `-l` option to list the information about the files that contain the partial path.

If you installed a previous Software Express release, your system might use a SQL format package database. The SQL database was created in one of the following ways.

- You performed an initial installation of a Solaris Express release prior to the Solaris Express 10/04 release.
- You upgraded to a previous Solaris Express release, and upgraded the package database manually by running `pkgadm upgrade`.

When you upgrade to the Solaris Express 10/04 release or subsequent releases, the SQL package database is automatically converted back to the ASCII text file format.

See the *System Administration Guide: Basic Administration* and the `pkgchk(1M)` man page for further information.

Enhancements for `pbind` and `psrset` Commands

These enhancements are new in the Solaris Express 10/04 release.

System administrators can now query and establish per-thread processor and processor set bindings. New options have been added to `pbind` and `psrset` commands. These new options enable you to query which LWPs on the system are bound to specified processor or processor sets. You can also clear all processor or processor set bindings.

See the `pbind(1M)` and `psrset(1M)` man pages for further information about these commands.

Solaris Dynamic Host Configuration Protocol (DHCP)

This feature is new in the Software Express pilot program. This feature is included in the Solaris 10 3/05 release.

The Solaris Dynamic Host Configuration Protocol (DHCP) service includes the DHCP Manager graphical tool and command-line tools. This Solaris release adds new options to the `dhcpconfig` and `dhtadm` commands.

With the new options for the `dhcpconfig` command, you can do the following:

- Enable and disable DHCP service for system starts
- Restart DHCP server
- Prompt DHCP server to read the `dhcptab`
- Specify DHCP server parameters

A new suboption for the `dhtadm` command causes the DHCP server to read the `dhcptab` to implement changes that were made with the `dhtadm` command.

The new options enable you to perform all DHCP management tasks from the command line as well as from DHCP Manager.

For more information, see the man pages for `dhcpconfig(1M)` and `dhtadm(1M)`.

See also the *System Administration Guide: IP Services*.

DHCP Event Scripting

This feature is new in the Solaris Express 8/04 release.

DHCP is a protocol for configuring systems with network information and leased IP addresses. The DHCP client is a daemon that runs on a network host and obtains configuration parameters for the host's network interfaces from a DHCP server.

The Solaris DHCP client daemon, `dhcpcagent`, is now able to execute administrator-defined actions. The actions can be triggered after any of the following DHCP lease events occur:

- The interface is configured for DHCP.
- The client successfully extends a lease.
- The lease expires.
- The client drops the lease to remove the interface from DHCP control.
- The client relinquishes the address.

The administrator-defined actions must be initiated by an executable script or program that you provide. You can use the program to perform any action that is appropriate for your system after one of the DHCP events.

For more information, see the [dhcpcagent\(1M\)](#) man page.

DHCP for Logical Interfaces

This feature is new in the Solaris Express 8/04 release.

The DHCP client daemon, `dhcpcagent`, can now be used with logical interfaces as well as physical interfaces. DHCP is started on logical interfaces by using the same methods that are used to start DHCP on physical interfaces.

However, the DHCP client has two limitations for the management of leases on logical interfaces:

- The DHCP client does not automatically manage default routes associated with logical interfaces.
- The DHCP client does not automatically generate a client identifier for logical interfaces.

For more information, see the [dhcpcagent\(1M\)](#) man page.

x86: SunVTS 6.0

SunVTS is a comprehensive diagnostic tool that tests and validates Sun hardware by verifying the connectivity and functionality of most hardware controllers and devices on Sun platforms.

SunVTS support for x86 systems is new in the Solaris Express 8/04 release and in the Solaris 10 3/05 release. The SunVTS infrastructure and a few core diagnostics are now available for x86 platforms. The current x86 support is for the 32-bit operating system only.

Note – The Solaris 10 1/06 release provides SunVTS 6.1. See [“SunVTS 6.1” on page 150](#).

You must install the x86 version of the SunVTS packages in order to use SunVTS on x86 platforms. The software packages use the same names as in the SPARC environment. A new package, `SUNWvtsr`, is provided for both the SPARC and x86 packages for this release. The

SUNWvtsr package contains the SunVTS framework configuration files in the root partition. The x86 packages are located in the `/sol_10_x86/s0/Solaris_10/ExtraValue/CoBundled/SunVTS_6.0` directory of the Solaris Software media. The SPARC packages are located in the `/sol_10_sparc/s0/Solaris_10/ExtraValue/CoBundled/SunVTS_6.0` directory of the Solaris Software media.

Refer to the SunVTS user documentation available at the following web sites for detailed information about x86 system support: <http://docs.sun.com> or <http://www.sun.com/products-n-solutions/hardware/docs/>.

Kernel Modular Debugger

Kernel Modular Debugger (KMDB) is an MDB-based kernel debugger. In the Solaris Express 8/04 release, KMDB replaced the existing kernel debugger (kadb) as the standard Solaris kernel debugger.

KMDB brings all the power and flexibility of Modular Debugger (MDB) to live kernel debugging. KMDB supports the following:

- `dcmds` – Debugger commands
- `dmods` – Debugger modules
- Access to kernel type data
- Kernel execution control
- Inspection
- Modification

KMDB can be loaded at boot time or after the system has already booted.

For more information, see the `kmdb(1M)` man page. See also the *Solaris Modular Debugger Guide* and the *System Administration Guide: Basic Administration*.

Solaris Process Accounting and Statistics Improvements

This feature is new in the Solaris Express 8/04 release.

In this release, changes have been made to the internals of the load averaging, `cpu usr/sys/idle`, and accounting functions. Microstate accounting has replaced the old accounting mechanism and is enabled by default. As a result, you might notice slightly different process usage and timing statistics.

Switching to microstate accounting provides substantially more accurate data about user processes and the amount of time they spend in various states. In addition, this information is used to generate more accurate load averages and statistics from the `/proc` file system.

For more information about process accounting and statistics, see the *System Administration Guide: Advanced Administration*. See also the `proc(4)` man page.

Enhancements to `ls` Command

This feature is new in the Solaris Express 8/04 release.

The `ls` command can now display file times to the second or to the nanosecond. Use the `-e` option for seconds and the `-E` option for nanoseconds. The following example illustrates these new options:

```
% ls -l foo
-rw-r--r--  1 jpk   staff   0 Aug  6 15:08 foo
% ls -e foo
-rw-r--r--  1 jpk   staff   0 Aug  6 15:08:28 2004 foo
% ls -E foo
-rw-r--r--  1 jpk   staff   0 2004-08-06 15:08:28.776641000 -0700 foo
```

See the [ls\(1\)](#) man page for further information about this command.

New `psrinfo` Option to Identify Chip Multithreading Features

This feature is new in the Solaris Express 6/04 release and in the Solaris 9 12/03 release.

In the Solaris 10 OS, the `psrinfo` command has been modified to provide information about physical processors, in addition to information about virtual processors. This enhanced functionality has been added to identify chip multithreading (CMT) features. The new `-p` option reports the total number of physical processors that are in a system. Using the `psrinfo -pv` command lists all the physical processors that are in the system, as well as the virtual processors that are associated with each physical processor.

The default output of the `psrinfo` command continues to display the virtual processor information for a system.

For more information, see the [psrinfo\(1M\)](#) man page. For information about the procedures that are associated with this feature, see the *System Administration Guide: Advanced Administration*.

Enhancements to `pfiles` Tool

This feature is new in the Solaris Express 6/04 release.

The `/proc` file system has been enhanced to include file name information in the `/proc/<pid>/path` directory. This information is used by [pfiles\(1\)](#) to display file names for each file in the process. This enhancement creates an entirely new form of observability that provides new insight into process behavior. This new observability benefits both administrators and developers.

Solaris IP Filter

This feature is new in the Solaris Express 2/04 release.

Solaris IP Filter is a host-based firewall that is based on the open source IP Filter. Solaris IP Filter provides packet filtering and network address translation (NAT), based on a user-configurable policy. Packet filtering rules are configurable in either a stateful or stateless manner. All configuration and management of Solaris IP Filter is performed through a command-line interface.

For further information, see the [ipfilter\(5\)](#) man page. Also, refer to the [ipf\(1M\)](#), [ipfs\(1M\)](#), and [ipfstat\(1M\)](#) man pages. In addition, see the *System Administration Guide: IP Services*.

Core File Content Enhancements

The Solaris Express 12/03 release introduced new enhancements to the `coreadm`, `gcore`, and `mdb` utilities. These changes improve management of core files as described in the following three sections. Further enhancements to the `coreadm` command were made in the Solaris Express 1/04 release.

Specify Core File Content With `coreadm`

This feature is new in the Solaris Express 12/03 release. The `-i` and `-I` options have been further enhanced in the Solaris Express 1/04 release.

In the Solaris Express 12/03 release, the `coreadm` command lets you specify which parts of a process are present in the core file during a crash. You can see the system's configuration by running `coreadm` with no arguments.

You can specify the global core file content and the default per-process core file content by using the `-G` and `-I` options respectively. Each option requires a set of content specifier tokens. You can also set the core file content for individual processes by using the `-P` option. Core dumps that correspond to the global settings no longer honor the per-process, core file-size resource control.

In the Solaris Express 1/04 release, the `-i` and `-I` options to the `coreadm` command now apply to all processes whose core file settings are using the system-wide default. Use the `-p` and `-P` options to override the default.

For further information, see the [coreadm\(1M\)](#) man page.

`gcore` Core File Content

This feature is new in the Solaris Express 12/03 release.

The `gcore` utility creates core files from a running process without damaging that process. The `gcore` utility now supports variable core file content. Use the `-c` option to specify the content or the `-p` or `-g` options to force `gcore` to use the `coreadm` settings.

See the [gcore\(1\)](#) man page for further information.

`mdb` Supports Text and Symbol Tables in Core Files

This feature is new in the Solaris Express 12/03 release.

Text is now in core files by default. Also, symbol tables can now be in core files by default. The `mdb` utility has been updated to support this new core file data. This support means you can now debug your old core file without needing the original binary or the libraries that are linked to that file.

See the [mdb\(1\)](#) man page for further information.

System Management Agent

This feature is new in the Solaris Express 12/03 release.

The System Management Agent is a Simple Network Management Protocol (SNMP) agent that provides SNMPv1, SNMPv2c, and SNMPv3 functionality to the Solaris 10 environment. The agent is based on the Net-SNMP open source implementation, with some customizations for the Solaris environment. The agent has all the necessary base functionality that is required by an SNMP agent. The agent includes support for standard SNMP operations and numerous standard Management Information Bases (MIBs), including MIB-II, Host Resources MIB, and Notification MIB. Additionally, the agent supports the User-based Security Model (USM) and View-based Access Control Model (VACM), as well as AgentX.

The System Management Agent is configured to be the default SNMP agent, but coexists with the Solstice Enterprise Agents™ software in this release.

For further information, see the [netsnmp\(5\)](#) man page.

For further information, see also the *Solaris System Management Agent Administration Guide* and the *Solaris System Management Agent Developer's Guide*.

SPARC: 64-bit Package Changes for System Administrators

This feature is new in the Software Express pilot program. This feature is included in the Solaris 10 3/05 release.

In Solaris software, packaging has been simplified with most 32-bit components and 64-bit components being delivered in a single package. For further information, see “[SPARC: 64-bit Package Changes](#)” on page 189.

NIS-to-LDAP Transition Service

This feature is new in the Software Express pilot program and in the Solaris 9 12/03 release. This feature is included in the Solaris 10 3/05 release.

The NIS-to-LDAP transition service can be used to enable a network transition from using NIS to using LDAP as the primary naming service. By using this transition service, administrators can utilize the Sun Java System Directory Server, which works with LDAP naming service clients.

For more information, see *System Administration Guide: Naming and Directory Services (DNS, NIS, and LDAP)*.

Signed Packages and Patches

This feature is new in the Software Express pilot program and in the Solaris 9 12/03 release. This feature is included in the Solaris 10 3/05 release.

The Solaris software enables you to securely download Solaris packages and patches that include a digital signature by using the updated `pkgadd` and `patchadd` commands. A package or a patch with a valid digital signature ensures that the package or patch has not been modified after the signature was applied to the package or patch.

In previous Solaris releases, you could only add signed patches to your system if you used the Solaris patch management tools with PatchPro 2.1.

Additional software management features in this Solaris release include the following:

- You can add a digital signature to a package with the updated `pkgtrans` command. For information about creating a signed package, see the [Application Packaging Developer's Guide](#).
- You can download a package or patch from an HTTP or HTTPS server.

A signed package is identical to an unsigned package except for the signature. The package can be installed, queried, or removed with existing Solaris packaging tools. A signed package is also binary-compatible with an unsigned package.

Before you can add a package or patch with digital signatures to your system, you must set up a keystore with trusted certificates that are used to identify that the digital signature on the package or patch is valid.

For information about setting up the package keystore and adding signed packages or patches to your system, see the [System Administration Guide: Basic Administration](#).

For information about booting and retrieving Solaris installation images from an HTTP or HTTPS server, see “WAN Boot Installation Method” on page 189.

System V IPC Configuration

This feature is new in the Software Express pilot program. This feature is included in the Solaris 10 3/05 release.

In the Solaris software, all System V IPC facilities are either automatically configured, or the facilities can be controlled by resource controls. These facilities might be shared memory, message queues, or semaphores. Resource controls allow IPC settings to be made on a per-project or per-user basis on the local system or in a name service environment. In previous Solaris releases, IPC facilities were controlled by kernel tunables. You had to modify the `/etc/system` file and reboot the system to change the default values for these facilities. Because the IPC facilities are now controlled by resource controls, their configuration can be modified while the system is running. Many applications that previously required system tuning to function might now run without tuning because of increased defaults and automatic allocation of resources.

The following table identifies the now obsolete IPC tunables that are described in the *Solaris Tunable Parameters Reference Manual* and their replacement resource controls.

Resource Control	Obsolete Tunable	Old Default Value	Maximum Value	New Default Value
process.max-msg-qbytes	msginfo_msgmnb	4096	ULONG_MAX	65536
process.max-msg-messages	msginfo_msgtql	40	UINT_MAX	8192
process.max-sem-ops	seminfo_semopm	10	INT_MAX	512
process.max-sem-nsems	seminfo_semmsl	25	SHRT_MAX	512
project.max-shm-memory	shminfo_shmmax	0x800000	UINT64_MAX	1/4 of physical memory
project.max-shm-ids	shminfo_shmmni	100	2 ²⁴	128
project.max-msg-ids	msginfo_msgmni	50	2 ²⁴	128
project.max-sem-ids	seminfo_semmni	10	2 ²⁴	128

Obsolete parameters might be included in the `/etc/system` file on this Solaris system. If so, the parameters are used to initialize the default resource control values as in previous Solaris releases. However, using the obsolete parameters is not recommended.

The following related parameters have been removed. If these parameters are included in the `/etc/system` file on this Solaris system, the parameters are commented out.

- `semsys:seminfo_semmns`
- `semsys:seminfo_sevmx`
- `semsys:seminfo_semmnu`
- `semsys:seminfo_semaem`
- `semsys:seminfo_semume`
- `semsys:seminfo_semusz`
- `semsys:seminfo_semmap`
- `shmsys:shminfo_shmseg`
- `shmsys:shminfo_shmmn`
- `msgsys:msginfo_msgmap`
- `msgsys:msginfo_msgseg`
- `msgsys:msginfo_msgssz`
- `msgsys:msginfo_msgmax`

For more information about using the new resource controls, see “[System V IPC and Other Resource Controls](#)” on page 180.

Optional Interval for `netstat` Command

This feature is new in the Software Express pilot program. This feature is included in the Solaris 10 3/05 release.

The `netstat` command displays the contents of certain network-related data structures in various formats, depending on the options that you select. The `-s`, `-i` and `-m` options now support an “interval.” With the specification of an optional interval, the output is continuously displayed in interval seconds until interrupted by the user. If a “count” is specified after “interval,” the output iterates for “count” times only. The “interval” and “count” must both be positive integers.

For further information, see the [netstat\(1M\)](#) man page.

GMT-offset Quoted Time Zones

This feature is new in the Software Express pilot program. This feature is included in the Solaris 10 3/05 release.

The Solaris Install program installs a quoted time zone as the system default time zone, if a GMT-offset style time zone is selected. Previously, an unquoted time zone was installed.

Note that quoted time zones behave differently than unquoted time zones. When the time zone name is displayed by the `date` command, for example, quoted time zones display the offset from GMT. Unquoted time zones display only the string “GMT,” without the offset from GMT.

Combine Data Address and Test Address for Single IP Network Multipathing Group

This feature is new in the Software Express pilot program and in the Solaris 9 12/02 release. This feature is included in the Solaris 10 3/05 release.

The use of a dedicated test IP address for failure detection on single-adapter IP network multipathing groups is no longer required. Sometimes, failover is not possible because only one network interface card (NIC) is in an IP network multipathing group. In this situation, you can now combine the test address and the data address. The `in.mpathd` daemon uses a data address to detect failure when a test address is not specified.

For further information on Solaris IP multipathing, see Part VI, “IPMP,” in the *System Administration Guide: IP Services*. See also the [in.mpathd\(1M\)](#) man page.

LDAP Command Changes

These command changes were made in the Solaris Express 12/03 release.

Several LDAP commands are updated to include full SSL support and extended support for SASL. The modifications also provide support for managing smart referrals, using virtual list views (VLVs), and establishing stronger authentication when binding to the LDAP server.

This upgrade aligns the Solaris LDAP command functionality with Sun's LDAP directory server commands. All Solaris functionality is preserved for backward compatibility. The updated commands are `ldapdelete`, `ldapmodify`, `ldapadd`, `ldapsearch`, and `ldapmodrdn`.

Several changes have been made to the LDAP commands. Those changes include the following:

- The `-M` authentication option is obsolete. This option has been superseded by the stronger `-o` option. The `-M` option is now used for managing smart referrals.
- Search results are now displayed in LDAP Data Interchange Format (LDIF) by default. Use the `-r` option to display results in the old format, for backward compatibility.

See the `ldapdelete(1)`, `ldapmodify(1)`, `ldapadd(1)`, `ldapsearch(1)`, and `ldapmodrdn(1)` man pages for details.

For further information, see the *System Administration Guide: Naming and Directory Services (DNS, NIS, and LDAP)*.

Desktop Enhancements

This section describes all desktop features in the Solaris 10 3/05 release that are new or have been enhanced since the Solaris 9 OS was originally distributed in May 2002. The Java Desktop System, Release 3 is an important addition to the Solaris 10 system. To view desktop features that are new in the Solaris 10 7/05 release, see “[Desktop Enhancements](#)” on page 153.

Java Desktop System, Release 3

This feature is new in the Solaris Express 10/04 release.

Note – Further enhancements are available in the Solaris 10 1/06 release. See in “[Enhancements to Java Desktop System Release 3](#)” on page 145.

The Sun Java Desktop System, Release 3 for the Solaris 10 OS combines open source software with Sun innovation to offer a comprehensive, fully integrated, and interoperable desktop environment with an accessible, modern user interface.

Comprehensive Environment

The Java Desktop System provides a comprehensive desktop environment that enables users to navigate and work effectively with the applications and documents on their desktop. The system includes the following:

- GNOME desktop environment
- StarOffice productivity suite
- Evolution mail, with calendar and tasks
- Mozilla browser and email client
- Complete set of software utilities and accessories

Intuitive User Interface

The Java Desktop System provides an intuitive desktop interface with familiar desktop themes and file manager views to enable users to work effectively and efficiently. Other features of the user interface include the following:

- Easy to learn and use, especially for users with Windows experience, featuring:
 - Familiar desktop themes for file management and desktop configuration
 - Standard keyboard shortcut conventions
 - Direct access to files and servers using the “Documents” and “Network Places” directories
 - Familiar access to networked computers from a single “Network Places” directory
- Comprehensive user documentation and Help

Integrated Desktop

Desktop applications are well integrated in the Java Desktop System, providing interoperability between applications. Interoperability features include the following:

- Drag-and-drop and copy or paste text, graphics, and other elements between applications
- Universal printer selection of available printers on the network, to include printers on Windows networks
- Single point of entry for directory and file management and networked computers
- Direct access to content in data and files by single-click application launch for files contained in e-mail, Web pages, and file manager

Enterprise-level Interoperability

In addition to providing interoperability between applications in the desktop, the Java Desktop System provides interoperability at the enterprise level. Features include the following:

- Users can access data, files, and printers with both Windows and UNIX environments.
- The system can connect to existing backend systems, including Sun Java System Calendaring and Messaging servers and Active Directory, POP3, LDAP, and SMTP mail servers.

For further information about the desktop, see the *Java Desktop System Release 3 User Guide*. For further information about the applications included in the Java Desktop System, open Help in the desktop and select the specific application from the navigation pane.

Accessibility Features

The Java Desktop System incorporates many new accessibility features to support people with disabilities. Features in the Solaris 10 OS include enhanced keyboard and desktop navigation,

which make the desktop accessible and customizable. The new assistive technologies included in the release are gnompernicus, a screen reader and screen magnifier, and a GNOME onscreen keyboard (GOK).

For further information, see the *Java Desktop System Release 3 Accessibility Guide*.

Mozilla 1.7

Mozilla 1.4 is new in the Solaris Express 4/04 release. Mozilla 1.7 is new in the Solaris Express 10/04 release.

Mozilla is Sun's new standard browser for the Solaris Operating System, delivering functionality that enables industry-standard web interaction.

Mozilla includes tools for the following:

- Web browsing
- Managing email
- Communicating with colleagues
- Participating in discussion groups
- Creating dynamic web pages

Mozilla includes the following new features:

- Non-Standard web page support
- Profile roaming support
- Junk-mail filtering
- Enhanced pop-up window blocking
- Easier bookmark management

For more information about Mozilla, go to <http://www.sun.com/software/solaris/browser/>.

System Resources Enhancements

This section describes all system resources enhancements in the Solaris 10 3/05 release that are new or have been enhanced since the Solaris 9 OS was originally distributed in May 2002. The Solaris 10 OS includes a feature of key importance to Solaris users, the Solaris Zones partitioning technology.

Solaris Zones Software Partitioning Technology

This feature is new in the Solaris Express 2/04 release. In the Solaris Express 7/04 release, new functionality for Zones has been added.

The Solaris Zones software partitioning technology, a component of the Solaris Containers environment, is a software partitioning technology used to virtualize operating system services and provide an isolated and secure environment for running applications. A zone is a

virtualized operating system environment created within a single instance of the Solaris Operating System. Zones basically provide the standard Solaris interfaces and application environment, and do not include a new ABI or API that would require applications to be ported.

Each zone can provide a customized set of services. Zones are ideal for environments that consolidate multiple applications on a single server. Resource management features can be used within zones to further control how applications use available system resources.

A zone can be thought of as a box. One or more applications can run in this box without affecting the rest of the system. This isolation prevents processes that are running in one zone from monitoring or interfering with processes that are running in other zones. Even a process with superuser credentials that is running inside a zone cannot view or affect activity in other zones.

The single instance of the Solaris Operating System is the global zone. The global zone is both the default zone for the system and the zone used for system-wide administrative control. One or more non-global zones can be created by an administrator working in the global zone. Once created, these non-global zones can be administered by individual zone administrators. The privileges of a zone administrator are confined to a non-global zone.

Non-global zones provide isolation at almost any level of granularity you require. A zone does not need a dedicated CPU, a physical device, or a portion of physical memory. These resources can either be multiplexed across several zones running within a single domain or system, or allocated on a per-zone basis using the resource management features available in the operating system. Even a small uniprocessor system can support multiple zones running simultaneously.

To achieve process isolation, a process can see or signal only those processes that exist in the same zone.

Basic communication between zones is provided by giving each zone at least one logical network interface. Applications running in different zones on the same system can bind to the same network port by using the distinct IP addresses associated with each zone or by using the wildcard address. An application running in one zone cannot observe the network traffic of another zone. This isolation is maintained even though the respective streams of packets travel through the same physical interface.

Each zone is given a portion of the file system hierarchy. Because each zone is confined to its subtree of the file system hierarchy, a workload running in a particular zone cannot access the on-disk data of another workload running in a different zone.

Files used by naming services reside within a zone's own root file system view. Thus, naming services in different zones are isolated from one other and can be configured differently.

For information about how to configure and use zones on your system, see the *System Administration Guide: Solaris Containers-Resource Management and Solaris Zones*.

System V IPC and Other Resource Controls

The resource controls (rctl) feature is new in the Solaris 9 software. Additional resource controls are new in the Software Express pilot program, and in the Solaris Express 3/04, 8/04, and 10/04 releases.

New project-based and new process-based resource controls have been added. Resource controls provide a mechanism for controlling how applications use system resources.

In the Software Express pilot program, the following resource controls were added. These resource controls affect the System V IPC (interprocess communication) facilities, which include shared memory, message queues, and semaphores.

- `project.max-shm-ids`
- `project.max-sem-ids`
- `project.max-msg-ids`
- `project.max-shm-memory`
- `process.max-sem-nsems`
- `process.max-sem-ops`
- `process.max-msg-qbytes`

See [“System V IPC Configuration” on page 173](#) for more information.

In the Solaris Express 3/04 release, the following event port resource controls were added:

- `project.max-device-locked-memory`
- `project.max-port-ids`
- `process.max-port-events`

In the Solaris Express 8/04 release, the following new resource controls have been added:

- `project.max-lwps`
- `project.max-tasks`

In the Solaris Express 10/04 release, the `project.max-contracts` resource control was added.

In the Solaris Express 11/04 release, the `process.max-crypto-memory` cryptographic resource control was added.

Information about resource controls is provided in Chapter 6, “Resource Controls (Overview),” in the *System Administration Guide: Solaris Containers-Resource Management and Solaris Zones*. This chapter includes the following:

- A current list of available resource controls
- A general description of resource controls
- Resource control configuration information

Information about configuring resource controls is also available in the `project(4)` man page.

New Solaris Project and Resource Management Command Functionality

This feature is new in the Solaris Express 7/04 release.

Enhancements to project database and resource control commands that are new in the Solaris Express 7/04 release include the following:

- Scaled value and unit modifier support for resource control values and commands
- Improved validation and easier manipulation of the project attributes field
See the [project\(4\)](#) man page.
- Revised output format and new options for the `prctl` and `projects` commands
See the [prctl\(1\)](#) and [projects\(1\)](#) man pages.

For information about the changes introduced, see the following chapters in *System Administration Guide: Solaris Containers-Resource Management and Solaris Zones*:

- Chapter 2, “Projects and Tasks (Overview)”
- Chapter 3, “Administering Projects and Tasks”
- Chapter 6, “Resource Controls (Overview)”
- Chapter 7, “Administering Resource Controls (Tasks)”

Note that information about setting the user's default project through the `useradd`, `usermod`, and `passmgmt` commands has been updated in the guide. These commands are documented in the [useradd\(1M\)](#), [usermod\(1M\)](#), and [passmgmt\(1M\)](#) man pages.

See also the following man pages:

- [prctl\(1\)](#)
- [projects\(1\)](#)
- [projadd\(1M\)](#)
- [projdel\(1M\)](#)
- [projmod\(1M\)](#)
- [rctladm\(1M\)](#)
- [setrctl\(2\)](#)
- [rctlblk_set_value\(3C\)](#)
- [setproject\(3PROJECT\)](#)
- [project\(4\)](#)

Dynamic Resource Pools

This feature is new in the Solaris Express 1/04 release.

Dynamic resource pools (DRPs) provide a mechanism for adjusting each pool's resource allocation in response to system events and application load changes on systems that have resource pools enabled. Adjustments are automatically made to maintain the system performance goals specified by an administrator. Changes that are made to the configuration are logged.

These features are primarily enacted through the resource controller `poold`, a system daemon that is active when dynamic resource allocation is required. Periodically, `poold` examines the load on the system and determines whether intervention is required to enable the system to maintain the specified resource utilization goals. The daemon takes corrective action if possible, or the condition is logged.

For further information, see the following chapters in the *System Administration Guide: Solaris Containers-Resource Management and Solaris Zones*:

- Chapter 12, “Dynamic Resource Pools (Overview)”
- Chapter 13, “Administering Dynamic Resource Pools (Tasks)”
- Chapter 14, “Resource Management Configuration Example”

See also the following man pages:

- `pooladm(1M)`
- `poolbind(1M)`
- `poolcfg(1M)`
- `poold(1M)`
- `poolstat(1M)`
- `libpool(3LIB)`

Extended Accounting Subsystem Enhancements

Extended accounting is new when the Solaris 9 Operating System was originally distributed in May 2002. The Software Express pilot program introduced flow accounting enhancements and a Perl interface. These enhancements are included in the Solaris 10 3/05 release.

Note – For Solaris 9 users, flow accounting enhancements are new in the Solaris 9 9/02 release. The Perl interface is new in the Solaris 9 4/03 release.

You can use extended accounting processes in conjunction with the flow accounting module for IPQoS. For information about IPQoS, see Part VII, “IP Quality of Service (IPQoS),” in the *System Administration Guide: IP Services*.

For information about the extended accounting feature, see Chapter 5, “Administering Extended Accounting (Tasks),” in the *System Administration Guide: Solaris Containers-Resource Management and Solaris Zones*.

A Perl interface to `libexacct` is available in the Software Express releases. This interface enables you to create Perl scripts that can read the accounting files which are produced by the `exacct` framework. You can also create Perl scripts that write `exacct` files. The new interface is functionally equivalent to the underlying C API.

You can use the Perl interface to record system resource consumption on a task or process basis. Or, you can record consumption on the basis of selectors provided by the IPQoS `flowacct` module.

For further information, see the following man pages:

- `Exacct(3PERL)`
- `Exacct::Catalog(3PERL)`
- `Exacct::File(3PERL)`
- `Exacct::Object(3PERL)`
- `Exacct::Object::Group(3PERL)`
- `Exacct::Object::Item(3PERL)`
- `Kstat(3PERL)`
- `Project(3PERL)`
- `Task(3PERL)`

For information on how to configure and use extended accounting with enhancements described in this section, see Chapter 4, “Extended Accounting (Overview),” in the *System Administration Guide: Solaris Containers-Resource Management and Solaris Zones*.

Physical Memory Control Using the Resource Capping Daemon

This feature is new in the Software Express pilot program and in the Solaris 9 12/03 release. This feature is included in the Solaris 10 3/05 release.

Physical memory control that uses the resource capping daemon is an optional feature. The resource capping daemon `rcapd` regulates the consumption of physical memory by processes that run in projects that have defined resource caps. Associated utilities provide mechanisms for administering the daemon and reporting related statistics.

For additional information, see Chapter 10, “Physical Memory Control Using the Resource Capping Daemon (Overview),” in the *System Administration Guide: Solaris Containers-Resource Management and Solaris Zones*.

See also the man pages `rcapstat(1)`, `rcapadm(1M)`, `project(4)`, and `rcapd(1M)`.

Enhancements to the Java 2 Platform, Standard Edition 5

This section describes Java platform features that are new in the Solaris Express 8/04 release.

Java 2 Platform, Standard Edition (J2SE™) 5 is a significant release including many new features and updates while preserving compatibility and stability. For more details about the features described in this section, see <http://java.sun.com/j2se/5.0>.

New Additions to the Java Programming Language

In the Solaris Express 8/04 release, the following features are now available to developers who use the Java language:

- The ability to declare type-safe generic types.
- Automatic conversion between primitive types and their object counterparts, referred to as “autoboxing” and “autounboxing.”
- The ability to create type-safe enumerated types, more specifically, an enhanced “for loop.”
- The ability to import a static namespace, an annotation API for use with packages, fields, methods, and classes.
- Permission to use subclassed return types, covariant returns, and variable argument lists. The javac compiler enables these features by default.

The Java programming language changes reduce the complexity and the amount of additional code that is required for common Java programming tasks. These changes also reduce the risk of runtime errors that a developer might introduce.

Monitoring and Manageability Enhancements for Java Platform

In the Solaris Express 8/04 release, the Java Virtual Machine (JVM) internals can now be monitored by industry-standard consoles that use JMX or SNMP protocols. Both mechanisms are disabled by default and are preconfigured to provide low memory notifications. The JMX mechanism also allows properties such as the logging level to be dynamically configured at runtime.

For further information, see <http://java.sun.com/j2se/1.5.0/docs/guide/management/index.html>.

Performance and Scalability Improvements for Java Platform

In the Solaris Express 8/04 release, J2SE 5 offers enhanced performance and scalability for developers who use the Java platform.

Automatic Tuning for Large Server Machines

Machines with more than 2 Gbytes of main memory and more than 2 processors automatically select an enhanced server performance mode. This mode delivers significant performance, especially for J2EE™ based application servers. This enhanced performance mode uses the “hotspot” server compiler with a parallel “garbage collector,” and with appropriately tuned initial and maximum memory sizes.

For further information, see <http://java.sun.com/j2se/1.5.0/docs/guide/vm/gc-ergonomics.html>.

Reduced Startup Time

Desktop applications benefit from the addition of class data sharing. During installation, this feature creates a special preprocessed archive of commonly used Java classes. This same archive is loaded when the JVM starts, thus reducing the processing time. This archive is also available to subsequent Java applications on that same machine.

OpenGL Hardware Acceleration

The JVM can accelerate JFC/Swing and Java 2D applications if an OpenGL device is available. This acceleration is automatic but disabled by default. This acceleration improves the runtime performance of many graphic-intensive programs.

XML 1.1 and Namespace, XSLTC, SAX 2.0.2, and DOM Level 3, Java based APIs

This feature is new in the Solaris Express 8/04 release.

Developers can now use the later XML 1.1 syntax. Developers can process XML by using the updates to the existing event-based XML processor, SAX 2.0.2. Developers can also use the newer document-based processor, DOM Level 3. For handling XML transforms, the fast XSLT compiler (XSLTC) is available.

New Default Swing Look and Feel

This feature is new in the Solaris Express 8/04 release.

The default cross-platform “look and feel” has been updated. Titled Ocean, this new “look and feel” is enabled by default. In addition, a GNOME-based GTK “look and feel” is now fully available for developers and users.

Diagnostic Tools

This feature is new in the Solaris Express 8/04 release.

In addition to using the monitoring protocols, administrators and developers can use new command-based diagnostic tools. Tools are available for the following tasks:

- Performance analysis – `jstat`
- Fatal error handling – `XX:OnError`
- Core file analysis – `jdi` connectors
- JVM internals tools – `jstack`, `jmap`, `jinfo`

In addition, Java extensions are available for the DTrace tracing facility.

Unicode Support

In the Solaris Express 8/04 release, Unicode 4.0-based supplementary characters are supported in J2SE 5.

Developers can now use Unicode-based supplementary characters, represented as pairs of Java characters, with the Java language. When using the string-based and text-based APIs, no additional effort is required.

For further information, see <http://java.sun.com/j2se/1.5.0/docs/>.

Note – Previously, Unicode 3.2 support was available in the Software Express pilot program. For Solaris 9 users, Unicode 3.2 support is new in the Solaris 9 8/03 release.

Installation Enhancements

This section describes all installation enhancements in the Solaris 10 3/05 release that are new or have been enhanced since the Solaris 9 OS was originally distributed in May 2002. The installation unification enhancements are of particular importance.

Solaris Installation Changes Including Installation Unification

Introduced in the Solaris Express 6/04 release, several changes in the installation of the Solaris OS provide a more simplified and unified installation experience.

The changes include the following:

- This release has one installation DVD and several installation CDs. The Solaris 10 Operating System DVD includes the content of all the installation CDs.
 - **Solaris Software 1** – This CD is the only bootable CD. From this CD, you can access both the Solaris installation graphical user interface (GUI) and the console-based installation. This CD also enables you to install selected software products from both the GUI and the console-based installation.
 - **Other Solaris Operating System CDs** – These CDs contain the following:
 - Solaris packages that the software prompts you to install if necessary
 - ExtraValue software that includes supported and unsupported software
 - Installers
 - Localized interface software and documentation
- The Solaris Installation CD no longer exists.
- For both CD and DVD media, the GUI installation is the default (if your system has enough memory). However, you can specify a console-based installation with the text boot option.
- The installation process has been simplified, enabling you to select the language support at boot time, but select locales later.

Note – The (noninteractive) Solaris custom JumpStart installation method has not changed.

To install the OS, simply insert the Solaris Software - 1 CD or the Solaris Operating System DVD and type one of the following commands.

- For the default GUI installation (if system memory permits), type **boot cdrom**.
- For the console-based installation, type **boot cdrom - text**.

For instructions about how to install the Solaris OS by using CD or DVD media with the new text boot option, see *Solaris 10 Installation Guide: Basic Installations*. For changes to setting up an installation server with CD media, see *Solaris 10 Installation Guide: Network-Based Installations*.

Accessing the GUI or Console-based Installations

You can choose to install the software with a GUI or with or without a windowing environment. Given sufficient memory, the GUI is displayed by default. If the memory is insufficient for the GUI, other environments are displayed by default. You can override defaults with the `nowin` or `text` boot options. But, you are limited by the amount of memory in your system or by installing remotely. Also, if the Solaris installation program does not detect a video adapter, the program is automatically displayed in a console-based environment. The following table describes these environments and lists minimal memory requirements for displaying them.

TABLE 9-1 Memory Requirements

Minimal Memory	Type of Installation
128–383 MB	Console-based
384 MB or greater	GUI-based

Detailed descriptions for each installation option are as follows:

Installation with 128–383 MB minimal memory

This option contains no graphics, but provides a window and the ability to open other windows. This option requires a local or remote DVD-ROM or CD-ROM drive or network connection, video adapter, keyboard, and monitor. If you install by using the `text` boot option and have enough memory, you are installing in a windowing environment. If you are installing remotely through a tip line or by using the `nowin` boot option, you are limited to the console-based installation.

Installation with 384 MB memory or greater

This option provides windows, pull-down menus, buttons, scrollbars, and iconic images. A GUI requires a local or remote DVD-ROM or CD-ROM drive or network connection, video adapter, keyboard, and monitor.

Custom JumpStart Installation Package and Patch Enhancements

This feature is new in the Solaris Express 3/04 release.

When installing and upgrading the Solaris OS by using the custom JumpStart installation method, new customizations enable the following:

- A Solaris Flash installation with additional packages
The custom JumpStart profile package keyword has been enhanced to enable installing a Solaris Flash archive with additional packages. For example, you can install the same base archive on two machines, but add a different set of packages to each machine. These packages do not have to be a part of the Solaris OS distribution.
- An installation with additional packages that might not be part of the Solaris distribution
The package keyword has also been enhanced to enable an installation with a package that is not part of the Solaris distribution. You no longer need to write a postinstallation script to add extra packages.
- An installation with the ability to install Solaris OS patches
The new custom JumpStart profile patch keyword enables the installation of Solaris OS patches. This feature allows the installation of a list of patches that are specified in a patch file.

For further information, see *Solaris 10 Installation Guide: Custom JumpStart and Advanced Installations*.

x86: Specifying Boot Properties With `add_install_client` Command

This feature is new in the Solaris Express 12/03 release and in the Solaris 9 4/04 release.

The `add_install_client` command now enables you to set boot properties for Solaris OS clients during a PXE network boot on x86 systems. The `-b` option enables you to perform the following tasks with the `add_install_client` command.

- You can specify an alternate console to use during your network installation.
- You can specify the device to use as the network boot device during the installation.
- You can instruct the client to perform a fully automated custom JumpStart installation.

For more information, see the `install_scripts(1M)` man page, the `eeeprom(1M)` man page, and *Solaris 10 Installation Guide: Network-Based Installations*.

Configuring Multiple Network Interfaces During Installation

This feature is new in the Solaris Express 12/03 release.

The Solaris installation programs enable you to configure multiple interfaces during your installation. You can preconfigure these interfaces in the `sysidcfg` file for your system. Or, you can configure multiple interfaces during the installation.

For more information, see the `sysidtool(1M)` man page and the `sysidcfg(4)` man page. See also the *Solaris 10 Installation Guide: Network-Based Installations*.

Replacement of Solaris Installation Command-Line Interface

This feature is new in the Solaris Express 9/03 release.

The command-line interface of the Solaris installation program has been replaced to improve the usability of console-based installations. All console-based installations use the forms-based program to install the Solaris Operating System.

See the [install-solaris\(1M\)](#) man page for further information.

SPARC: 64-bit Package Changes

This feature is new in the Software Express pilot program. This feature is included in the Solaris 10 3/05 release.

In previous Solaris releases, the Solaris software was delivered in separate packages for 32-bit components and 64-bit components. Now, packaging has been simplified with most 32-bit components and 64-bit components being delivered in a single package. The combined packages retain the names of the original 32-bit packages, and the 64-bit packages are no longer delivered.

The removal of the 64-bit packages simplifies installation and increases performance:

- Reduces the number of packages, which simplifies Custom JumpStart scripts that contain lists of packages
- Simplifies the packaging system with only one package that groups software functions
- Reduces installation time because fewer packages are installed

The 64-bit packages are renamed with the following conventions:

- If a 64-bit package has a 32-bit counterpart, the 64-bit package is named with the 32-bit package name. For example, a 64-bit library such as `/usr/lib/sparcv9/libc.so.1` previously would have been delivered in `SUNWcs1x`, but now is delivered in `SUNWcs1`. The 64-bit `SUNWcs1x` package is no longer delivered.
- If a package does not have a 32-bit counterpart, the “x” suffix is removed from the name. For example, `SUNW1394x` becomes `SUNW1394`.

This change means that you might need to modify your custom JumpStart script or other package installation scripts to remove references to the 64-bit packages.

WAN Boot Installation Method

This feature is new in the Software Express pilot program and in the Solaris 9 12/03 release. This feature is included in the Solaris 10 3/05 release.

The Solaris software now enables you to boot and install software over a wide area network (WAN) by using HTTP. The WAN boot installation method enables you to install the Solaris

release on systems over a large public network where the network infrastructure might be untrustworthy. You can use WAN boot with new security features to protect data confidentiality and installation image integrity.

The WAN boot installation method enables you to transmit an encrypted Solaris Flash archive over a public network to a remote client. The WAN boot programs then install the client system by performing a custom JumpStart installation. To protect the integrity of the installation, you can use private keys to authenticate and encrypt data. You can also transmit your installation data and files over a secure HTTP connection by configuring your systems to use digital certificates.

For more information about this feature, see the [Solaris 10 Installation Guide: Network-Based Installations](#).

Creating RAID-1 Volumes (Mirrors) With Custom JumpStart

This feature is new in the Software Express pilot program and in the Solaris 9 4/04 release. This feature is included in the Solaris 10 3/05 release.

The custom JumpStart installation method now enables you to create RAID-1 volumes (mirrors) during the installation of the Solaris Operating System. By mirroring file systems, you can protect your system by duplicating data over two physical disks. If one of the mirrored disks fails, the system data is still accessible from the second mirrored disk.

In JumpStart, the following new custom profile keywords and values enable you to create mirrored file systems.

- The new `filesys` keyword value `mirror` creates a mirror. Then you can designate specific slices as single-slice concatenations to attach to the mirror.
- The new `metadb` profile keyword enables you to create the required state database replicas.

For more information, see [Solaris 10 Installation Guide: Custom JumpStart and Advanced Installations](#).

Solaris Live Upgrade 2.1

This feature is new in the Software Express pilot program and in the Solaris 9 8/03 release. This feature is included in the Solaris 10 3/05 release.

Solaris Live Upgrade provides a method of upgrading a system while the system continues to operate. While your current boot environment is running, you can duplicate the boot environment, then upgrade the duplicate. Or, rather than upgrading, you can install a Solaris Flash archive on a boot environment. The original system configuration remains fully functional and unaffected by the upgrade or installation of an archive. When you are ready, you can activate the new boot environment by rebooting the system. If a failure occurs, you have a safety net. You can quickly revert to the original boot environment with a simple reboot. Thus, you eliminate the normal downtime of the test and evaluation process.

New features in Solaris Live Upgrade 2.1 provide the following new functionality.

- Solaris Live Upgrade uses Solaris Volume Manager technology to create a duplicate boot environment that contains file systems with RAID-1 volumes (mirrors). The mirror provides data redundancy for any file systems, including the root (/) file system. With the `lucreate` command, you can create mirrored file systems that contain up to three submirrors.
- With the `lucreate` command, you can now exclude some files and directories that would normally be copied from the original boot environment. If you have excluded a directory, you can also include specified files and subdirectories under that directory.

For further information, see *Solaris 10 Installation Guide: Solaris Live Upgrade and Upgrade Planning*.

Custom JumpStart Installation Method Creates New Boot Environment

This feature is new in the Software Express pilot program. This feature is included in the Solaris 10 3/05 release.

You can now use the JumpStart installation method to create an empty boot environment when you install the Solaris Operating System. The empty boot environment can then be populated with a Solaris Flash archive for later use.

For further information, see *Solaris 10 Installation Guide: Custom JumpStart and Advanced Installations*.

Solaris Flash Archives

The Solaris Flash features that are described in this section are new in the Software Express pilot program. This feature is included in the Solaris 10 3/05 release.

The Solaris Flash installation feature enables you to use a single reference installation of the Solaris software on a system. This system is called the master system. Then, you can replicate that installation on multiple systems, which are called clone systems. The installation is an initial installation that overwrites all files on the clone system.

Solaris Flash Differential Archives and Configuration Scripts

This feature is new in the Software Express pilot program and in the Solaris 9 4/03 release. This feature is included in the Solaris 10 3/05 release.

The Solaris Flash installation feature provides new enhancements for the Solaris software.

- A Solaris Flash installation can now update a clone system with minor changes. If you have a clone system and want to update it, you can create a differential archive that contains only the differences between two images, the original master image and an updated master image. When you update a clone system with a differential archive, only the files that are specified in the differential archive are changed. The installation is restricted to clone systems that contain software which is consistent with the original master image. You use the custom JumpStart installation method to install a differential archive on a clone system. Or, you can use Solaris Live Upgrade to install a differential archive on a duplicate boot environment.
- Special scripts can now be run for configuration of the master or clone or can be run to validate the archive. These scripts enable you to do the following tasks.
 - Configure applications on clone systems. You can use a custom JumpStart script for some uncomplicated configurations. For more complicated configurations, special configuration file processing might be necessary on the master system or before or after installation on the clone system. Also, local preinstallation and postinstallation scripts can reside on the clone. These scripts can protect local customizations from being overwritten by the Solaris Flash software.
 - Identify nonclonable, host-dependent data that enables you to make the flash archive host independent. Host independence is enabled by modifying such data or excluding the data from the archive. An example of host-dependent data is a log file.
 - Validate software integrity in the archive during creation.
 - Validate the installation on the clone system.

For further information, see *Solaris 10 Installation Guide: Solaris Flash Archives (Creation and Installation)*. This guide also includes information on how to use Solaris Live Upgrade to install a differential archive.

Customizing Content in a Solaris Flash Archive

This feature is new in the Software Express pilot program and in the Solaris 9 12/02 release. This feature is included in the Solaris 10 3/05 release.

The `flarcreate` command is used to create a Solaris Flash archive. This command has been updated with new options that increase your flexibility to define archive contents when creating an archive. You now can exclude more than one file or directory. From an excluded directory, you can add back a subdirectory or file. This feature is useful when you want to exclude large data files that you do not want cloned.

For information about how to use these options, see *Solaris 10 Installation Guide: Solaris Flash Archives (Creation and Installation)*.

Note – Solaris Flash was formerly Web Start Flash.

Command-Line Interface Enhancements to the Solaris Product Registry

These enhancements are new in the Software Express pilot program and in the Solaris 9 4/03 release. This feature is included in the Solaris 10 3/05 release.

The `prodreg` command has been updated to include functionality that is similar to the Solaris Product Registry graphical user interface. You can now use the following `prodreg` subcommands on the command line or in administration scripts to perform a variety of tasks.

- `browse` – The `browse` subcommand enables you to view registered software in a terminal window. By repeating the `browse` subcommand, you can navigate through the directory hierarchy of registered software.
- `info` – The `info` subcommand enables you to view information about registered software.

You can use the `info` subcommand to identify the following specifics.

- Installation location of the software
- Other software that the specified software requires
- Other software that depends on the specified software
- Software that has been damaged by the removal of packages that the software requires
- `unregister` – The `unregister` subcommand removes software installation information from the Solaris Product Registry. If you remove software from your system without properly uninstalling the software from the Registry, you can use the `prodreg unregister` command to clean up the obsolete entries in the Solaris Product Registry.
- `uninstall` – The `uninstall` subcommand enables you to remove registered software from your system by launching that software's uninstallation program.

For more information, see the [prodreg\(1M\)](#) man page and the *System Administration Guide: Basic Administration*.

Reduced Networking Software Group

This feature is new in the Software Express pilot program. This feature is included in the Solaris 10 3/05 release.

You can now create a more secure system with fewer enabled network services by selecting or specifying the Reduced Networking software group (SUNWCrnet) during your installation. The Reduced Networking software group provides system administration utilities and a multiuser text-based console. SUNWCrnet enables the system to recognize network interfaces. During the installation, you can customize your system's configuration by adding software packages and by activating network services as needed.

For further information, see *Solaris 10 Installation Guide: Custom JumpStart and Advanced Installations*.

Modifying Disk Partition Tables by Using a Virtual Table of Contents (VTOC)

This feature is new in the Software Express pilot program. This feature is included in the Solaris 10 3/05 release.

The Solaris installation program now enables you to load existing slices from the virtual table of contents (VTOC.) You can now preserve and use the system's existing disk slice tables during your installation, rather than using the installer's default disk layout.

x86: Change in Default Boot-Disk Partition Layout

This feature is new in the Software Express pilot program. This feature is included in the Solaris 10 3/05 release.

A new feature in the Solaris installation program is a boot-disk partition layout. This layout, by default, accommodates the Service partition on Sun x86 based systems. This installation program enables you to preserve an existing Service partition.

The new default includes the following partitions.

- First partition – Service partition (existing size on system)
- Second partition – x86 boot partition (approximately 11 Mbytes)
- Third partition – Solaris Operating System partition (remaining space on the boot disk)

If you want to use this default layout, select Default when the Solaris installation program asks you to choose a boot-disk layout.

Note – If you install the Solaris OS for x86 based systems on a system that does not currently include a Service partition, the Solaris installation program no longer automatically creates a new Service partition. If you want to create a Service partition on your system, first use your system's diagnostic CD to create the Service partition. After the Service partition is created, then install the Solaris Operating System.

For information about how to create the Service partition, see your hardware documentation.

For more information, see *Solaris 10 Installation Guide: Custom JumpStart and Advanced Installations*.

Support for Lightweight Directory Access Protocol (LDAP) Version 2 Profiles

This feature is new in the Software Express pilot program and in the Solaris 9 12/02 release. This feature is included in the Solaris 10 3/05 release.

The Solaris installation programs now support LDAP Version 2 profiles. These profiles enable you to configure your system to use a proxy credential level. During the Solaris installation programs, you can specify the LDAP proxy-bind distinguished name and proxy-bind password. With any installation method, you can preconfigure LDAP before installation by using the `proxy_dn` and `proxy_password` keywords in the `sysidcfg` file.

For information, see *Solaris 10 Installation Guide: Network-Based Installations*.

Security Enhancements

This section describes security enhancements in the Solaris 10 3/05 release that are new or have been enhanced since the Solaris 9 OS was originally distributed in May 2002. “[Process Rights Management](#)” on page 195 and “[Reduced Networking Software Group](#)” on page 193 are of particular importance. To view security enhancements that are new in the Solaris 10 7/05 release, see “[Security Enhancements](#)” on page 147.

In addition to the security features described in this section, see also the following security-related feature descriptions in the Developer Tools section and the Installation section:

- “[WAN Boot Installation Method](#)” on page 189
- “[Reduced Networking Software Group](#)” on page 193
- “[New Mechanisms for the Solaris Cryptographic Framework](#)” on page 237
- “[Retail and Nonretail Options for Providers in Solaris Cryptographic Framework](#)” on page 237
- “[Simple Authentication and Security Layer for Developers](#)” on page 240
- “[SPNEGO Pseudo-Mechanism for GSS-API Applications](#)” on page 244
- “[Enhanced crypt \(\) Function for Software Developers](#)” on page 245
- “[Smart Card Terminal Interfaces](#)” on page 246
- “[Smartcard Middleware APIs](#)” on page 246

Signing ELF Objects

This feature is new in the Solaris 10 3/05 release.

The libraries and executable files in the Solaris 10 OS include digital signatures that can be used to verify the integrity of these files. The digital signature provides a way to detect any accidental change or deliberate tampering with the executable content of the file.

Plug-ins for the Solaris Cryptographic Framework are automatically verified when loaded by the system. The `elfsign` command can be used manually to verify any signed file. Developers and administrators can also use `elfsign` to sign their own code.

For further information, see the `elfsign(1)` man page.

Process Rights Management

This feature is new in the Software Express pilot program. This feature is included in the Solaris 10 3/05 release.

In the Solaris software, administrative tasks that previously required superuser rights are now protected by process rights management. Process rights management uses privileges to restrict processes at the command, user, role, or system level. A privilege is a discrete right that a process requires to perform an operation. The system restricts processes to only those privileges that are required to perform the current task. Therefore, fewer root processes are vulnerable to exploitation. The number of `setuid` programs has been greatly reduced.

As installed, the Software Express releases and the Solaris 10 3/05 release are completely compatible with previous releases of the Solaris Operating System in terms of the privileges enhancements. Unmodified programs that run as root run with all privileges.

Device Protection – Devices are protected with a security policy. The policy is enforced with privileges. Therefore, the permissions on a device file do not fully determine the device's availability. Privileges might also be required to operate the device.

System interfaces that were protected by UNIX permissions are now protected by privileges. For example, members of the group `sys` are no longer automatically allowed to open the `/dev/ip` device. Processes that are running with the `net_rawaccess` privilege can access the `/dev/ip` device. When the system boots, access to all devices is restricted until the `devfsadm` command runs during the boot sequence. The initial policy is as strict as possible. The policy prevents all users except the superuser from initiating connections.

See the following man pages for more information:

- [getdevpolicy\(1M\)](#)
- [ppriv\(1\)](#)
- [add_drv\(1M\)](#)
- [update_drv\(1M\)](#)
- [rem_drv\(1M\)](#)
- [devfsadm\(1M\)](#)

Processes that need to retrieve Solaris IP MIB information should open `/dev/arp` and push the “tcp” and “udp” modules. No privileges are required. This method is equivalent to opening `/dev/ip` and pushing the “arp”, “tcp” and “udp” modules. Because opening `/dev/ip` now requires a privilege, the `/dev/arp` method is preferred.

For further information, see the following sections in the *System Administration Guide: Security Services*:

- “Using Roles and Privileges (Overview)”
- “Privileges (Overview)”
- “Privileges (Tasks)”

Changes to PAM for the Solaris 10 OS

A new `pam_deny` module was added in the Software Express pilot program and enhanced in the Solaris Express 6/04 release. This feature is included in the Solaris 10 3/05 release. The module can be used to deny access to named PAM services. By default, the `pam_deny` module is not used. For more information, see the [pam_deny\(5\)](#) man page.

The Solaris 10 software includes the following changes to the PAM framework.

- The `pam_authok_check` module now allows for strict password checking that uses new tunables in the `/etc/default/passwd` file. The new tunables define the following:
 - A list of comma-separated dictionary files that are used for checking common dictionary words in a password

- The minimum differences that are required between a new password and an old password
- The minimum number of alphabetic and nonalphabetic characters that must be used in a new password
- The minimum number of uppercase and lowercase letters that must be used in a new password
- The number of allowable consecutive repeating characters
- The number of digits that must be used in the new password
- Whether whitespaces are allowed in the new password
- The `pam_unix_auth` module implements account locking for local users. Account locking is enabled by the `LOCK_AFTER_RETRIES` tunable in `/etc/security/policy.conf` and the `lock_after-retries` key in `/etc/user_attr`.
- A new binding control flag has been defined. If the PAM module is successful and no preceding modules that are flagged as required have failed, then PAM skips the remaining modules and the authentication request succeeds. However, if a failure is returned, PAM records a required failure and then continues processing the stack. This control flag is documented in the [pam.conf\(4\)](#) man page.
- The `pam_unix` module has been removed and replaced by a set of service modules of equivalent or greater functionality. Many of these modules are new in the Solaris 9 system. Here is a list of the replacement modules:
 - `pam_authtok_check`
 - `pam_authtok_get`
 - `pam_authtok_store`
 - `pam_dhkeys`
 - `pam_passwd_auth`
 - `pam_unix_account`
 - `pam_unix_auth`
 - `pam_unix_cred`
 - `pam_unix_session`
- The functionality of the `pam_unix_auth` module has been split into two modules. The `pam_unix_auth` module now verifies that the password is correct for the user. The new `pam_unix_cred` module provides functions that establish user credential information.
- Additions to the `pam_krb5` module have been made to manage the Kerberos credentials cache by using the PAM framework. See [“Kerberos Enhancements” on page 200](#).

`pam_ldap` Changes

The following `pam_ldap` changes are new in the Solaris Express 10/04 release, except for the account management feature. This management feature is new in the Software Express pilot program and in the Solaris 9 12/02 release. See the [pam_ldap\(5\)](#) man page for more information about these changes.

- The previously supported `use_first_pass` and `try_first_pass` options are obsolete as of this Solaris 10 software release. These options are no longer needed. The options can safely be removed from `pam.conf`, and are silently ignored.
- Password prompting must be provided for by stacking `pam_authtok_get` before `pam_ldap` in the authentication and password module stacks, and by including `pam_passwd_auth` in the `passwd_service_auth` stack.
- The previously supported password update function is replaced in this release by the use of `pam_authtok_store` with the `server_policy` option.
- The `pam_ldap` account management feature strengthens the overall security of the LDAP Naming Service. Specifically, the account management feature does the following:
 - Allows for tracking password aging and expiration
 - Prevents users from choosing trivial or previously used passwords
 - Warns users if their passwords are about to expire
 - Locks out users after repeated login failures
 - Prevents users other than the authorized system administrator from deactivating initialized accounts

Note – A clean, automated update cannot be provided for the changes in the previous list. Therefore, an upgrade to a Solaris 10 or subsequent release cannot automatically update the existing `pam.conf` file to reflect the `pam_ldap` changes. If the existing `pam.conf` file contains a `pam_ldap` configuration, the CLEANUP file notifies you after the upgrade. Examine the `pam.conf` file and modify it, as needed.

See the following man pages for further information:

- `pam_passwd_auth(5)`
- `pam_authtok_get(5)`
- `pam_authtok_store(5)`
- `pam.conf(4)`

For further information about Solaris naming and directory services, see the [System Administration Guide: Naming and Directory Services \(DNS, NIS, and LDAP\)](#). For information about Solaris security features, see the [System Administration Guide: Security Services](#).

Solaris Secure Shell Enhancements

This feature description is new in the Solaris Express 10/04 release.

The following enhancements to Solaris Secure Shell are available in the Solaris 10 OS:

- Solaris Secure Shell is based on OpenSSH 3.5p1. The Solaris implementation also includes features and bug fixes from versions prior to OpenSSH 3.8p1.

- Solaris Secure Shell now supports the use of the GSS-API for user and host authentication that uses Kerberos V.
PAM support, including password-aging support, has been improved.
- The default value of `X11Forwarding` is `yes` in the `/etc/ssh/sshd_config` file.
- The ARCFOUR and AES128-CTR ciphers are now available. ARCFOUR is also known as RC4. The AES cipher is AES in counter mode.
- For further enhancements, see the description for “[sshd Daemon and /etc/default/login](#)” on page 199.

For further information about security in the Solaris 10 OS, see the *System Administration Guide: Security Services*.

OpenSSL and OpenSSL PKCS#11 Engine

This feature is new in the Solaris Express 8/04 release.

This release of Solaris includes the OpenSSL libraries and commands in `/usr/sfw`.

This release also includes an OpenSSL Engine interface to PKCS#11 so that OpenSSL consumers can access hardware and software cryptographic providers from the Solaris Cryptographic Framework.

Note – Because of cryptographic import restrictions in some countries, symmetric key cryptographic algorithms are limited to 128-bit if the `SUNWcry` package is not installed. The `SUNWcry` package is not included with the Solaris software. This package is available instead as a separate controlled download.

sshd Daemon and `/etc/default/login`

This feature is new in the Solaris Express 10/04 release.

The `sshd` daemon uses the variables in `/etc/default/login` and the `login` command. The `etc/default/login` variables can be overridden by values in the `sshd_config` file.

For more information, see “Solaris Secure Shell and Login Environment Variables” in the *System Administration Guide: Security Services*. See also the `sshd_config(4)` man page.

New Password Options for Nonlogin and Locked Accounts

This feature is new in the Solaris Express 10/04 release.

The `passwd` command has two new options, `-N` and `-u`. The `-N` option creates a password entry for a nonlogin account. This option is useful for accounts that should not be logged in to, but must run `cron` jobs. The `-u` option unlocks a previously locked account.

For more information, see the `passwd(1)` man page.

-setcond **Option to auditconfig Command Is Removed**

This feature is new in the Solaris Express 10/04 release.

The `-setcond` option to the `auditconfig` command has been removed. To temporarily disable auditing, use the `audit -t` command. To restart auditing, use the `audit -s` command.

perzone **Audit Policy**

This feature is new in the Solaris Express 8/04 release.

The `perzone` audit policy enables non-global zones to be audited individually. A separate audit daemon runs in each zone. The daemon uses audit configuration files that are specific to the zone. Also, the audit queue is specific to the zone. By default, the policy is off.

For more information, see the [auditd\(1M\)](#) and [auditconfig\(1M\)](#) man pages.

Kerberos Enhancements

These Kerberos enhancements are included in the Solaris 10 release. Several of the enhancements are new in prior Software Express releases.

- Kerberos protocol support is provided in remote applications, such as `ftp`, `rcp`, `rdist`, `rlogin`, `rsh`, and `telnet`. See the man pages for each command or daemon and the [krb5_auth_rules\(5\)](#) man page for more information.
- The Kerberos principal database can now be transferred by incremental update instead of transferring the entire database each time. Incremental propagation provides several advantages including the following:
 - Increased database consistencies across servers
 - The need for fewer resources, such as network and CPU resources
 - Much more timely propagation of updates
 - An automated method of propagation
- A new script helps automatically configure a Kerberos client. The script helps an administrator quickly and easily set up a Kerberos client. For procedures that use the new script, see Chapter 22, “Configuring the Kerberos Service (Tasks),” in the *System Administration Guide: Security Services*. See also the [kclient\(1M\)](#) man page for more information.
- Several new encryption types have been added to the Kerberos service. These new encryption types increase security and enhance compatibility with other Kerberos implementations that support these encryption types. All of the encryption types are documented in the [mech\(4\)](#) man page. See “Using Kerberos Encryption Types” in the *System Administration Guide: Security Services* for more information. The encryption types offer the following capabilities:
 - AES encryption type can be used for high-speed, high-security encryption of Kerberos sessions. The use of AES is enabled through the Cryptographic Framework.
 - ARCFOUR-HMAC provides better compatibility with other Kerberos versions.

- Triple DES (3DES) with SHA1 increases security. This encryption type also enhances interoperability with other Kerberos implementations that support this encryption type.
- The KDC software and `kinit` command now support the use of the TCP network protocol. This addition provides more robust operation and better interoperability with other Kerberos implementations. The KDC now “listens” on both the traditional UDP ports and on the TCP ports so that it can respond to requests that use either protocol. The `kinit` command first tries UDP when sending a request to the KDC. If a failure occurs, the `kinit` command then tries TCP.
- Support for IPv6 was added to the KDC software with `kinit`, `klist`, and `kprop` commands. Support for IPv6 addresses is provided by default. No configuration parameters need to change to enable this support.
- A new `-e` option has been added to several subcommands of the `kadmin` command. This new option allows for the selection of the encryption type when creating principals. See the [kadmin\(1M\)](#) man page for more information.
- Additions to the `pam_krb5` module manage the Kerberos credentials cache by using the PAM framework. See the [pam_krb5\(5\)](#) man page for more information.
- Support is provided for auto-discovery of the Kerberos KDC, admin server, `kpaswd` server, and host or domain name-to-realm mappings that use DNS lookups. This support reduces some of the steps that are needed to install a Kerberos client. The client is able to locate a KDC server by using DNS instead of reading a configuration file. See the [krb5.conf\(4\)](#) man page for more information.
- A new PAM module called `pam_krb5_migrate` has been introduced. The new module helps in the automatic migration of users to the local Kerberos realm if the users do not already have Kerberos accounts. See the [pam_krb5_migrate\(5\)](#) man page for more information.
- The `~/ .k5login` file can now be used with GSS applications, `ftp` and `ssh`. For more information, see the [krb5_auth_rules\(5\)](#) man page.
- The `kproplog` utility has been updated to display all attribute names per log entry. For more information, see the [kproplog\(1M\)](#) man page.
- A new configuration file option makes the strict Ticket Granting Ticket (TGT) verification feature optionally configurable on a per-realm basis. See the [krb5.conf\(4\)](#) man page for more information.
- Extensions to the password-changing utilities enable the Solaris Kerberos V5 administration server to accept password change requests from non-Solaris clients. See the [kadmin\(1M\)](#) man page for more information.
- The default location of the replay cache has been moved from RAM-based file systems to persistent storage in `/var/krb5/rcache/`. The new location protects against replays if a system is rebooted. Performance enhancements were made to the `rcache` code. However, overall replay cache performance might be slower because of the use of persistent storage.
- The replay cache can now be configured to use file storage or memory-only storage. Refer to the [krb5envvar\(5\)](#) man page for more information about environment variables that can be configured for key table and credential cache types or locations.

- The GSS credential table is no longer necessary for the Kerberos GSS mechanism. For more information, see the [gsscred\(1M\)](#), [gssd\(1M\)](#), and [gsscred.conf\(4\)](#) man pages.
- The Kerberos utilities, `kinit` and `ktutil`, are now based on MIT Kerberos version 1.2.1. This change added new options to the `kinit` command and new subcommands to the `ktutil` command. For more information, see the [kinit\(1\)](#) and the [ktutil\(1\)](#) man pages.
- The Solaris Kerberos Key Distribution Center (KDC) is now based on MIT Kerberos version 1.2.1. The KDC now defaults to a btree-based database, which is more reliable than the current hash-based database. See the [kdb5_util\(1M\)](#) man page for more information. For Solaris 9 users, this change is new in the Solaris 9 12/03 release.

TCP Wrappers for `rpcbind`

This feature is new in the Solaris Express 4/04 release.

TCP wrapper support has been added to the `rpcbind` command. This support allows the administrator to limit calls to `rpcbind` to selected hosts. The administrator can also log all calls to `rpcbind`.

For further information, see the [rpcbind\(1M\)](#) man page.

zonename Audit Token and Audit Policy Option

The Solaris Zones partitioning technology is new in the Solaris Express 2/04 release. See “[Solaris Zones Software Partitioning Technology](#)” on page 178. The related `zonename` enhancements that are described here were also introduced in the Solaris Express 2/04 release.

The `zonename` audit token records the name of the zone in which an audit event occurred. The `zonename` audit policy option determines, for all zones, whether the `zonename` token is included in audit records. If the criteria for audit class preselection vary between non-global zones, then you might want to analyze audit records by zone. The `zonename` audit policy enables you to postselect audit records by zone.

See “Auditing and Solaris Zones” in the *System Administration Guide: Security Services*.

For further information, see the [audit.log\(4\)](#), [auditconfig\(1M\)](#), and [auditreduce\(1M\)](#) man pages. See also “Using Solaris Auditing in Zones” in the *System Administration Guide: Solaris Containers-Resource Management and Solaris Zones*.

User Commands for Solaris Cryptographic Framework

This feature is new in the Solaris Express 1/04 release.

The `digest`, `mac`, and `encrypt` commands now include an option to list the algorithms that are available for each command. For the `mac` and `encrypt` commands, the output includes the key lengths that each algorithm accepts. Also, the `-I <IV-file>` option has been removed from the `encrypt` and `decrypt` commands.

For further information, see Chapter 14, “Solaris Cryptographic Framework (Tasks)” and “Protecting Files With the Solaris Cryptographic Framework” in the *System Administration Guide: Security Services*.

For more information, see the `encrypt(1)`, `digest(1)`, and `mac(1)` man pages.

IKE Configuration Parameters

This feature is new in the Solaris Express 1/04 release.

Retransmission parameters and packet time-out parameters have been added to the `/etc/inet/ike/config` file. The parameters enable the administrator to tune the IKE Phase 1 (Main Mode) negotiation. The tuning enables Solaris IKE to interoperate with platforms that implement the IKE protocol differently. The parameters also help the administrator adjust for network interference and heavy network traffic.

For a detailed description of the parameters, see the `ike.config(4)` man page.

Simple Authentication and Security Layer

This feature is new in the Solaris Express 12/03 release.

Simple Authentication and Security Layer (SASL) provides developers of applications with interfaces for adding authentication, data integrity checking, and encryption to connection-based protocols.

For further information, see “Simple Authentication and Security Layer for Developers” on page 240.

See also Chapter 17, “Using SASL,” in the *System Administration Guide: Security Services*.

Audit Time Now Reported in ISO 8601 Format

This feature is new in the Solaris Express 12/03 release.

The file and header tokens in audit records now report time in ISO 8601 format. For example, the output from the `praudit` command for the file token is as follows:

Old File Token:

```
file,Mon Oct 13 11:21:35 PDT 2003, + 506 msec,  
/var/audit/20031013175058.20031013182135.machine1
```

New File Token:

```
file,2003-10-13 11:21:35.506 -07:00,  
/var/audit/20031013175058.20031013182135.machine1
```

Old Header Token:

```
header,173,2,setppriv(2),,machine1,  
Mon Oct 13 11:23:31 PDT 2003, + 50 msec
```

New Header Token:

```
header,173,2,setppriv(2),,machine1,  
2003-10-13 11:23:31.050 -07:00
```

The XML output has also changed. For example, the output from the `praudit -x` command formats the file token as follows:

```
<file iso8601="2003-10-13 11:21:35.506 -07:00">  
/var/audit/20031013175058.20031013182135.machine1</file>
```

Customized scripts or tools that parse `praudit` output might need to be updated to accommodate this change.

For further information, see Chapter 27, “Solaris Auditing (Overview)” and “Changes to Solaris Auditing for the Solaris 10 Release” in the *System Administration Guide: Security Services*.

Basic Audit and Reporting Tool

This feature is new in the Solaris Express 11/03 release.

Basic Audit and Reporting Tool (BART) is a command-line utility that enables OEMs, advanced users, and system administrators to do a check at file level of the software contents of a target system. The utility is useful for gathering information about what is installed on a system. BART also enables you to compare installed systems, and to compare the contents of a system over time.

For further information, see Chapter 5, “Using the Basic Audit Reporting Tool (Tasks),” in the *System Administration Guide: Security Services*.

See also the `bart_manifest(4)`, `bart_rules(4)`, and `bart(1M)` man pages.

IPsec and the Solaris Cryptographic Framework

This feature is new in the Solaris Express 9/03 release.

IPsec uses the Solaris Cryptographic Framework instead of its own encryption and authentication modules. The modules are optimized for the SPARC platform. In addition, a new `ipsecalgs` command-line utility and APIs are provided to query the list of supported IPsec algorithms and other IPsec properties.

For further information, see the `ipsecalgs(1M)` man page.

In the *System Administration Guide: IP Services*, see Chapter 18, “IP Security Architecture (Overview)” and “Authentication and Encryption Algorithms in IPsec.”

Solaris Cryptographic Framework for System Administrators

This feature is new in the Software Express pilot program. This feature is included in the Solaris 10 3/05 release.

The Solaris Cryptographic Framework provides cryptographic services to applications in the Solaris environment. The system administrator controls which encryption algorithms can be used through the `cryptoadm` command. The `cryptoadm` command enables you to perform the following functions:

- Manage available providers of cryptographic services
- Set cryptographic security policy, such as disabling algorithms from a particular provider

The framework comes with plug-ins for the AES, DES/3DES, RC4, MD5, SHA-1, DSA, RSA, and Diffie-Hellman algorithms. Plug-ins can be added or removed as needed.

The `encrypt`, `decrypt`, `digest`, and `mac` commands all use cryptographic algorithms from the framework.

For further information, see Chapter 13, “Solaris Cryptographic Framework (Overview),” in the *System Administration Guide: Security Services*.

See also the following man pages:

- `cryptoadm(1M)`
- `kcf(1M)`
- `libpkcs11(3LIB)`
- `pkcs11_kernel(5)`
- `pkcs11_softtoken(5)`

Remote Audit Log

This feature is new in the Software Express pilot program. This feature is included in the Solaris 10 3/05 release.

In addition to recording audit events in the binary audit log, the Solaris releases enable you to record audit events to `syslog`.

The generation of `syslog` data allows you to use the same management and analysis tools that are available for `syslog` messages from a variety of Solaris and non-Solaris environments, including workstation, servers, firewalls, and routers. By using `syslog.conf` to route audit messages to remote storage, you protect log data from alteration or deletion by an attacker.

However, the `syslog` option provides only a summary of audit record data. Also, when `syslog` data is stored on a remote system, the data is susceptible to network attacks such as denial of service and false or “spoofed” source addresses.

For further information, see Chapter 27, “Solaris Auditing (Overview)” and “Audit Files” in the *System Administration Guide: Security Services*.

See also the following man pages:

- `audit(1M)`
- `audit.log(4)`
- `audit_control(4)`
- `audit_syslog(5)`
- `syslog(3C)`
- `syslog.conf(4)`

FTP Server Enhancements

This feature is new in the Software Express pilot program. This feature is included in the Solaris 10 3/05 release.

Scalability and transfer logging enhancements have been made to the FTP server including:

- The `sendfile()` function is used for binary downloads.
- New capabilities are supported in the `ftppaccess` file:
 - `flush-wait` controls the behavior at the end of a download or directory listing.
 - `ipcos` sets the IP Class of Service for either the control or data connection.
 - `passive ports` can be configured so that the kernel selects the TCP port to listen on.
 - `quota-info` enables retrieval of quota information.
 - `recvbuf` sets the receive (upload) buffer size that is used for binary transfers.
 - `rhostlookup` allows or disallows the lookup of the remote host's name.
 - `sendbuf` sets the send (download) buffer size that is used for binary transfers.
 - `xferlog format` customizes the format of the transfer log entry.
- A new `-4` option makes the FTP server only listen for connections on an IPv4 socket when running in standalone mode.

The FTP client and server now support Kerberos. For more information, refer to the `ftp(4)` man page and to “Kerberos User Commands” in the *System Administration Guide: Security Services*.

In addition, `ftpcount` and `ftpwho` now support the `-v` option, which displays user counts and process information for FTP server classes that are defined in virtual host `ftppaccess` files.

For further information about these changes, see the following man pages:

- [in.ftpd\(1M\)](#)
- [ftpaccess\(4\)](#)
- [ftpcount\(1\)](#)
- [ftpwho\(1\)](#)
- [sendfile\(3EXT\)](#)

FTP Client

This feature is new in the Software Express pilot program. This feature is included in the Solaris 10 3/05 release.

The Solaris software includes a change to the FTP client. By default, a Solaris FTP client, connected to a Solaris FTP server, lists directories as well as plain files when the `ls` command is issued. If the FTP server is not running in the Solaris operating system, directories may not be listed.

To allow for the default Solaris behavior when connecting to non-Solaris FTP servers, the `/etc/default/ftp` file can be edited on each Solaris client. To make the change for individual users, the `FTP_LS_SENDS_NLST` environment variable can be set.

For more information, see the [ftp\(4\)](#) man page.

The FTP client and server now support Kerberos. For more information, refer to the [ftp\(4\)](#) man page and to “Kerberos User Commands” in the *System Administration Guide: Security Services*.

Internet Key Exchange (IKE) Key Storage on Sun Crypto Accelerator 4000 Board

This feature is new in the Software Express pilot program and in the Solaris 9 12/03 release. This feature is included in the Solaris 10 3/05 release.

IKE now runs on IPv6 as well as IPv4 networks. For information about keywords that are specific to the IPv6 implementation, see the [ifconfig\(1M\)](#) and [ike.config\(4\)](#) man pages.

When a Sun Crypto Accelerator 4000 board is attached, IKE can offload computation-intensive operations to the board, thus freeing the operating system for other tasks. IKE can also use the attached board to store public keys, private keys, and public certificates. Key storage on a separate piece of hardware provides additional security.

For further information, see the [ikecert\(1M\)](#) man page.

See also the following in *System Administration Guide: IP Services*:

- “IP Security Architecture (Overview)”
- “Internet Key Exchange (Overview)”

- “IKE and Hardware Storage”
- “Configuring IKE (Tasks)”
- “Configuring IKE to Find Attached Hardware”

IKE Hardware Acceleration

This feature is new in the Software Express pilot program and in the Solaris 9 4/03 release. This feature is included in the Solaris 10 3/05 release.

Public-key operations in IKE can be accelerated by a Sun Crypto Accelerator 1000 card and Sun Crypto Accelerator 4000 cards. The operations are offloaded to the card. The offloading accelerates encryption and reduces demands on Solaris Operating System resources.

For information about IKE, see the following in *System Administration Guide: IP Services*:

- “Configuring IKE to Find Attached Hardware”
- “Internet Key Exchange (Overview)”
- “IKE and Hardware Acceleration”
- “Configuring IKE (Tasks)”
- “Configuring IKE to Find Attached Hardware”

ipseckey Enhancement

This feature is new in the Software Express pilot program. This feature is included in the Solaris 10 3/05 release.

For network administrators who have installed IPsec or IKE on their systems, the `ipseckey` parser provides clearer help. The `ipseckey monitor` command now provides a timestamp for each event.

For more information, see the `ipseckey(1M)` man page.

Credential Propagation Over Loopback Connections

This feature is new in the Software Express pilot program. This feature is included in the Solaris 10 3/05 release.

This Solaris release introduces `ucred_t *` as an abstract representation of the credentials of a process. These credentials can be retrieved by using `door_ucred()` in door servers and `getpeerucred()` for loopback connections. The credentials can be received by using `recvmmsg()`.

See the `socket.h(3HEAD)` man page for further information.

Auditing Header Token Contains Host Information

This feature is new in the Software Express pilot program. This feature is included in the Solaris 10 3/05 release.

The header token in audit records has been expanded to include the name of the host.

The old header is displayed as follows:

```
header,131,4,login - local,,Wed Dec 11 14:23:54 2002, + 471 msec
```

The new expanded header is displayed as follows:

```
header,162,4,login - local,,example-hostname,
Fri Mar 07 22:27:49 2003, + 770 msec
```

Customized scripts or tools that parse `praudit` output might need to be updated to reflect this change.

See Chapter 30, “Solaris Auditing (Reference)” and “header Token” in the *System Administration Guide: Security Services* for further information.

Auditing Enhancements

This feature is new in the Software Express pilot program and in the Solaris 9 8/03 release. This feature is included in the Solaris 10 3/05 release.

Enhancements to the audit features in the Solaris software reduce noise in the trail, and enable administrators to use XML scripting to parse the trail. These enhancements include the following:

- Public files are no longer audited for read-only events. The `public` policy flag for the `auditconfig` command controls whether public files are audited. By not auditing public objects, the audit trail is greatly reduced. Attempts to read sensitive files are therefore easier to monitor.
- The `praudit` command has an additional output format, XML. The XML format enables the output to be read in a browser, and provides source for XML scripting for reports. See the `praudit(1M)` man page.
- The default set of audit classes has been restructured. Audit metaclasses provide support for more specific audit classes. See the `audit_class(4)` man page.
- The `bsmconv` command no longer disables the use of the Stop-A key combination. The Stop-A event is now audited to maintain security.

For further information, see the following sections in the *System Administration Guide: Security Services*:

- “Solaris Auditing (Reference)”
- “Definitions of Audit Classes”
- “praudit Command”
- “Solaris Auditing (Overview)”
- “Audit Terminology and Concepts”
- “Changes to Solaris Auditing for the Solaris 10 Release”

New Audit Token `path_attr`

This feature is new in the Software Express pilot program. This feature is included in the Solaris 10 3/05 release.

The `path_attr` audit token contains access path information for an object. The access path specifies the sequence of attribute file objects below the path token object. Systems calls such as `openat()` access attribute files. For more information on extended file attributes, see the [fsattr\(5\)](#) man page.

The `path_attr` token has three fields:

- A token ID field that identifies this token as a `path_attr` token
- A count that represents the number of sections of attribute file paths
- One or more null-terminated strings

The `praudit` command displays the `path_attr` token as follows:

```
path_attr,1,attr_file_name
```

For further information, see Chapter 30, “Solaris Auditing (Reference)” and “`path_attr` Token” in the *System Administration Guide: Security Services*.

Password History Checking

This feature is new in the Software Express pilot program. This feature is included in the Solaris 10 3/05 release.

For login accounts that are defined in local files, a password history of up to 26 previously changed passwords can be enabled. When a user changes a password, the attempt fails if the new password matches one of the passwords in the history. Also, login name checking can be disabled.

See the [passwd\(1\)](#) man page for more information.

Enhanced `crypt()` Function

This feature is new in the Solaris Express pilot program and in the Solaris 9 12/02 release. This feature is included in the Solaris 10 3/05 release.

Password encryption protects passwords from being read by intruders. Three strong password encryption modules are now available in the software:

- A version of Blowfish that is compatible with Berkeley Software Distribution (BSD) systems
- A version of Message Digest 5 (MD5) that is compatible with BSD and Linux systems
- A stronger version of MD5 that is compatible with other Solaris systems

For information on how to protect your user passwords with these new encryption modules, see the following sections in the *System Administration Guide: Security Services*:

- “Controlling Access to Systems (Tasks)”
- “Managing Machine Security (Overview)”
- “Changing the Default Algorithm for Password Encryption”

For information on the strength of the modules, see the `crypt_bsdbf(5)`, `crypt_bsmd5(5)`, and `crypt_sunmd5(5)` man pages.

System Performance Enhancements

This section describes all system performance features in the Solaris 10 3/05 release that are new or have been enhanced since the Solaris 9 OS was originally distributed in May 2002. The “New Architecture for Network Stacks” description is of particular importance.

New Architecture for Network Stacks

This new architecture is of key importance in the Solaris 10 system. This feature is new in the Solaris Express 10/03 release.

The networking stack for TCP connections has been rearchitected to deliver extremely high performance while improving the scalability. This innovation enables Solaris networking to vertically partition the workload by using an IP classifier-based lockless design. This design reduces the overheads of synchronization and reduces cross-communication between CPUs. The improvements are visible across all networking workloads and are transparent to the users.

CPU Performance Counters

The feature is new in the Solaris Express 4/04 release.

The CPU Performance Counter (CPC) system gives better access to the performance analysis features available on SPARC and x86 processors.

The CPC commands, `cpustat` and `cpurack`, have enhanced command-line syntax for specifying CPU information. In previous versions of the Solaris OS, you were required to specify two counters. The configuration of both commands now enables you to specify only one counter. For simple measurements, you can even omit the counter configuration.

For more information, see the [cpustat\(1M\)](#) man page. For more information about using the `cpurack` command, see the [cpurack\(1\)](#) man page.

Improved System Performance for Large Numbers of Interfaces

This feature is new in the Software Express pilot program. This feature is included in the Solaris 10 3/05 release.

System performance during and after plumbing large numbers of interfaces, such as virtual local area networks (VLANs) and tunnels, has been improved significantly. However, avoid using large numbers (thousands) of routes with large numbers (thousands) of interfaces.

UFS Logging Performance Enhancements

This feature is new in the Software Express pilot program and in the Solaris 9 12/02 release. This feature is included in the Solaris 10 3/05 release.

UFS logging enables faster system reboots. Because the file system transactions are already stored, checking the file system is not required if the file system is already consistent.

In addition, the performance of UFS logging improves or exceeds the level of performance of nonlogging file systems in this Solaris release. Results of the Standard Performance Evaluation Corporation system file server (SPECsfs) benchmark show that the performance of NFS-mounted file systems that have logging enabled equals the level of non-UFS logging performance. In some I/O-bound configurations, logging UFS file systems exceed the performance of nonlogging UFS file systems by approximately 25 percent. In other tests, performance of logging UFS file systems is 12 times better than the performance of nonlogging UFS file systems.

For information on how to enable logging on a UFS file system, see the [System Administration Guide: Devices and File Systems](#) or the [mount_ufs\(1M\)](#) man page.

Memory Placement Optimization (MPO)

This feature is new in the Software Express pilot program and in the Solaris 9 9/02 release. This feature is included in the Solaris 10 3/05 release.

The Solaris dispatcher and the virtual memory subsystem are enhanced to optimize the time that is required for an application to access memory. This enhancement can automatically

improve the performance of many applications. This feature is currently implemented to support specific platforms such as the Sun Fire 3800–6800, the Sun Fire 12K, and the Sun Fire 15K.

A technical white paper about *Solaris Memory Placement Optimization and Sun Fire Servers* is available at the following Web site:

<http://sundoc.central.sun.com/dm/index.jsp>

Dynamic Intimate Shared Memory (DISM) Large-Page Support

This feature is new in the Software Express pilot program and in the Solaris 9 9/02 release. This feature is included in the Solaris 10 3/05 release.

Large-page support is now provided for Dynamic Intimate Shared Memory (DISM). This large-page support can increase the performance of applications that are capable of dynamically adjusting shared memory sizes.

For further information on shared memory operations, see [shmop\(2\)](#).

Device Management

This section describes all device management features in the Solaris 10 3/05 release that are new or have been enhanced since the Solaris 9 OS was originally distributed in May 2002. To view device management features that are new in the Solaris 10 7/05 release, see “[Device Management Enhancements](#)” on page 155.

New or Updated Drivers

These driver features were new or revised in the Solaris Express 10/04 release.

bge Gigabit Ethernet Driver for Broadcom bcm57xx

The Solaris Express 10/04 release provided an enhanced bge driver to support Broadcom 5700/5701/5705/5782 gigabit Ethernet chipsets.

For further information, see the [bge\(7D\)](#) man page.

asy Handles IXON/IXOFF in Driver

This feature enables a quick response for software flow control on x86 platforms. This quick response avoids situations where asy can stall because an XOFF or XON character is embedded in pending data.

For further information, see the [asy\(7D\)](#) man page.

ECP Parallel Port Driver

The lp driver that was originally used for parallel print ports on x86 platforms lacked Enterprise CRM Platform (ECP) mode support. The new ECPP driver provides a robust parallel port driver with ECP mode support for x86 systems.

For further information, see the [ecpp\(7D\)](#) man page.

UHCI Driver Enhancement

An interrupt-out transfer has been added to the new Universal Host Controller Interface (UHCI) driver for x86 platforms. This new feature enables users to transfer urgent data to the devices through Interrupt Out in a high priority.

Other new features provided by the new UHCI driver include the following:

- BULK short transfer
- Port enumeration
- Multiple BULK transfer
- Isoc IN transfer
- BULK IN transfer

In the Solaris 10 3/05 release, the UHCI driver is available for the both x86 and SPARC platforms.

For further information, see the [uhci\(7D\)](#) man page.

Adaptec AdvancedRAID Controllers Support

Software support of Adaptec AdvancedRAID controllers has been added for x86 platforms. This feature enables you to utilize SCSI RAID disk arrays both as data drives and as boot drives. SCSI CD-ROM, DVD, or tapes are not supported, nor are online configurations.

For further information, see the [aac\(7D\)](#) man page.

Dell/LSI SCSI RAID PERC Controllers Support

Software support of Dell/LSI SCSI RAID PERC controllers has been added for x86 platforms. This feature enables users to utilize SCSI RAID disk arrays both as data drives and as boot drives. SCSI CD-ROM, DVD, or tapes are not supported, nor are online configurations.

For further information, see the [amr\(7D\)](#) man page.

Realtek RTL8139 NIC Support

This Solaris 10 OS provides x86 platform support for the Realtek fast Ethernet chipset RTL8139.

For further information, see the [rtls\(7D\)](#) man page.

audio810 Audio Driver

The `audio810` device driver has been added for x86 platforms. This driver supports the audio controllers that are embedded in the Intel ICH3, ICH4, ICH5, and AMD 8111 southbridge chips. This driver supports Solaris audio mixer interfaces, and enables users to utilize audio hardware to play back or record sound.

For further information, see the [audio810\(7D\)](#) man page.

SATA Controllers Support

Software support of selected SATA controllers was added for x86 platforms. This feature allows users to utilize Serial ATA Hard Drives both as data drives and as boot drives.

For further information, see the [ata\(7D\)](#) man page.

1394 (FireWire) and Mass-Storage Support on x86 Systems

This feature is new in the Solaris Express 10/04 release.

In this Solaris release, the 1394 OpenHCI host controller driver has been updated to include support for x86 systems. Previously support for 1394 (FireWire) technology was only available on SPARC systems.

IEEE 1394 is also known by the Apple Computer trademark name, FireWire.

1394 is an industry-standard serial bus that supports data rates of 100 Mbit/sec, 200 Mbit/sec, or 400 Mbit/sec. The bus readily handles data from consumer electronics devices, such as video cameras, because of its high bandwidth and isochronous (on-time) capabilities.

For more information, see the [hci1394\(7D\)](#) man page.

In this Solaris release, the `scsa1394` driver has been introduced to support 1394 mass-storage devices that are compliant with the Serial Bus Protocol 2 (SBP-2) specification. This driver supports both bus-powered and self-powered 1394 mass storage devices. Previously, only 1394 video cameras were supported.

1394 mass storage devices are treated as removable media devices. A 1394 mass storage device can be formatted by using the `rmformat` command. Using a 1394 mass-storage device is no different than using a USB mass-storage device. You can mount, eject, hot-remove, and hot-insert a 1394 mass-storage device.

For more information about using these devices, see the [hci1394\(7D\)](#) man page. See also Chapter 8, “Using USB Devices (Tasks),” in the *System Administration Guide: Devices and File Systems*.

Internet Printing Protocol Listener

This feature is new in the Solaris Express 10/04 release.

The Internet Printing Protocol (IPP) listener feature for printing on a network works in conjunction with the Apache web server that is shipped as part of the Solaris software. IPP uses HTTP to transport requests. Using the Apache web server, the listener loads IPP print requests and processes them. The IPP listener listens on port 631 for HTTP requests and then communicates the requests to the printing system.

For more information, see the *System Administration Guide: Advanced Administration*.

Fibre Channel Connectivity for Storage Devices

This feature is new in the Solaris Express 8/04 release.

The Solaris Fibre Channel I/O stack provides connectivity and multipathing functions for Fibre Channel devices in a Storage Area Network (SAN). The following enhancements have been made to this feature:

- Boot support with Solaris installation programs on SPARC platforms for Fibre Channel devices
- x86 platform support (32-bit and 64-bit) for Fibre Channel devices

Note – Previous SAN releases were distributed through SunSolveSM and Sun Download Center (SDLC) in patch sets and packages for the Solaris 8 OS and the Solaris 9 OS. All new SAN features will be integrated and distributed through Solaris releases.

For more information about SAN and related documents, refer to <http://www.sun.com/storage/san/>.

Expanded Printer Support

This feature is new in the Solaris Express 6/04 release and in the Solaris 9 9/04 release.

In this release, modifications have been made to incorporate support for a wide array of printers. This support is accomplished through the use of additional transformation software, raster image processor (RIP), and PostScriptTM Printer Description (PPD) files.

These additions provide functionality that enables you to print to printers, such as the Lexmark Optra E312 and Epson Stylus Photo 1280, by using PPD files.

This feature is useful in an environment where printers do not have resident PostScript processing capabilities.

In addition, the existing Solaris printing tools have been modified to include a new `-n` option to the `lpadmin` command. With this option, you can designate a PPD file to use when creating a new print queue or when modifying an existing print queue.

Also, the Solaris Print Manager screens have been updated to enable you to choose a PPD file for the print queue through the selection of make, model, and driver. This new feature differs greatly from previous Solaris software releases. In previous releases, the provided list of printer types, and information about whether the printer accepted PostScript or ASCII text, was limited.

For more information, see the `lpadmin(1M)` man page and the *System Administration Guide: Advanced Administration*.

Common Solaris Target Disk Driver

This feature is new in the Solaris Express 5/04 release.

In this Solaris release, the disk drivers for SPARC and x86 platforms are merged into a single driver. This change creates one source file for the following 3 drivers:

- SPARC `sd` for SCSI devices
- x86 `sd` for Fibre Channel and SCSI devices
- SPARC `ssd` for Fibre Channel devices

In previous Solaris releases, three separate drivers were needed to provide support for SCSI and Fibre Channel disk devices on the SPARC and x86 platforms. All of the disk utilities, such as the `format`, `fmthard`, and `fdisk` commands, have been updated to support these changes.

For more information, see the `sd(7D)` and `ssd(7D)` man pages.

In addition, Solaris support for the EFI disk label is now available on x86 systems. For further information about EFI disk labels, see “Multiterabyte Disk Support With EFI Disk Labels” on page 250.

For further information, see also “What’s New in Disk Management in the Solaris 10 Release?” in the *System Administration Guide: Devices and File Systems*.

Wheel Mouse Support

Support for the following mouse features was added in the Solaris Express 4/04 release and in the Solaris 9 9/04 release:

- Wheel mouse scrolling is available on a USB mouse device

This support means that rolling the wheel on a USB or a PS/2 mouse results in a “scroll” in the application or window under mouse focus.

StarOffice and Mozilla applications support wheel mouse scrolling. However, other applications might not support wheel mouse scrolling.

- Support for more than 3 buttons on USB or PS/2 mouse devices

USB 2.0 Features

This description is new in the Solaris Express 2/03 release and updated in the Solaris Express 9/03 release. For Solaris 9 users, USB 2.0 support is new in the Solaris 9 4/04 release.

The Software Express releases include the following USB 2.0 features:

- **Better Performance** – Increased data throughput for devices attached to USB 2.0 controllers, up to 40 times faster than USB 1.1 devices
You can particularly benefit from the high-speed USB protocol when accessing high-speed mass storage devices, such as DVDs and hard drives.
- **Backward Compatibility** – Compatibility with 1.0 and 1.1 devices and drivers so that you can use the same connectors and software interfaces
USB 1.x devices can also use the same 1.x cables as before or the newer USB 2.0 cables. USB 2.0 devices, per next description, require USB 2.0 cables.

For a description of USB devices and terminology, see “Overview of USB Devices” in the *System Administration Guide: Devices and File Systems*.

USB 2.0 Devices

This description is new in the Solaris Express 2/03 release and updated in the Solaris Express 9/03 release. For Solaris 9 users, USB 2.0 device support is new in the Solaris 9 4/04 release.

USB 2.0 devices are defined as high-speed devices that follow the USB 2.0 specification. You can refer to the USB 2.0 specification at <http://www.usb.org>.

Some of the USB devices that are supported in this Solaris release are as follows:

- Mass storage devices – CD-RWs, hard disks, DVD, digital cameras, Zip drives, diskettes, and tape drives
- Keyboard, mouse devices, speakers, and microphones
- Audio devices

For a full listing of USB devices that have been verified on the Solaris release, go to:

http://www.sun.com/io_technologies/ihvindex.html

Additional storage devices might work by modifying the `scsa2usb.conf` file. For more information, see the `scsa2usb(7D)` man page.

Solaris USB 2.0 device support includes the following features:

- Increased USB bus speed from 12 Mbps to 480 Mbps. This increase means devices that support the USB 2.0 specification can run significantly faster than their USB 1.1 counterparts when they are connected to a USB 2.0 port.

A USB 2.0 port is defined on SPARC and x86 systems as follows:

- A port on a USB 2.0 PCI card
- A port on a USB 2.0 hub that is connected to a USB 2.0 port
- A system USB port, assuming that it has a USB 2.0 port on the motherboard
- USB 2.0 is Solaris Ready on all PCI-based platforms. A USB 2.0 PCI card is needed to provide USB 2.0 ports. For a list of USB 2.0 PCI cards that have been verified for the Solaris release, go to http://www.sun.com/io_technologies/ihvindex.html.
- USB 1.1 devices work as they have in the past, even if you have both USB 1.1 and USB 2.0 devices on the same system.
- While USB 2.0 devices operate on a USB 1.x port, their performance is significantly better when connected to a USB 2.0 port.

For more information on USB 2.0 device support, see the [ehci\(7D\)](#) and [usba\(7D\)](#) man pages.

For information about USB cables and bus-powered devices, see “About USB in the Solaris OS” in the *System Administration Guide: Devices and File Systems*.

Solaris Support for USB Devices

This description is new in the Solaris Express 9/03 release.

USB 2.0 support for x86 systems was first introduced in the Solaris 9 4/04 release.

Support of USB 2.0 devices in the Solaris 8 and Solaris 9 releases is provided through the USBA 1.0 version of the USB Dual Framework. The USBA 2.0 framework in the Solaris 10 3/05 release supersedes the USBA 1.0 version.

USB Mass Storage Devices

This description is new in the Solaris Express 2/03 release and updated in the Solaris Express 9/03 release. For Solaris 9 users, this feature is new in the Solaris 9 4/04 release.

All USB storage devices in the Software Express releases are now accessed as removable media devices through volume management. This change has the following advantages:

- USB storage devices with standard MS-DOS or Windows (FAT) file systems are now supported.
- You can use the user-friendly `rmformat` command instead of the `format` command to format and partition all USB storage devices. If you need the functionality of the `format` command, use the `format -e` command.
- You can use the `fdisk` command if you need to do `fdisk`-style partitioning.
- Nonroot users can now access USB storage devices, because the root-privileged `mount` command is no longer needed. The device is automatically mounted by `vol` and is available under the `/rmdisk` directory. If a new device is connected while the system is down, do a reconfiguration boot with the `boot -r` command so that `vol` recognizes the device. If a new device is connected while the system is up, restart `vol`. For more information, refer to the [vol\(1M\)](#) and [scsa2usb\(7D\)](#) man pages.
- Disks with FAT file systems can be mounted and accessed. See the following example.

```
mount -F pcfs /dev/dsk/c2t0d0s0:c /mnt
```

- All USB storage devices are now power managed, except for those that support LOG SENSE pages. Devices with LOG SENSE pages are usually SCSI drives that are connected through a USB-to-SCSI bridge device. In previous Solaris releases, some USB storage devices were not power managed because they were not recognized as removable media.
- Applications might work differently with USB mass storage devices.

Note the following issues when using applications with USB storage devices.

- Applications might make incorrect assumptions about the size of the media because only smaller devices such as diskettes and Zip drives were removable previously.
- Requests by applications to eject media on devices where this would be inapplicable, such as a hard drive, succeed and do nothing.

To revert to the behavior of previous Solaris releases that did not treat all USB mass storage as removable media devices, update the `/kernel/drv/scsa2usb.conf` file.

For more information about using USB mass storage devices, see the [scsa2usb\(7D\)](#) man page.

For information on troubleshooting USB mass storage device problems, see “What’s New in USB Devices?” in the *System Administration Guide: Devices and File Systems*.

USB Driver Enhancements

This description is new in the Solaris Express 2/03 release and updated in the Solaris Express 9/03 release. For Solaris 9 users, this feature is new in the Solaris 9 4/04 release.

This section describes USB driver enhancements in the Software Express releases.

- **New generic USB driver** – USB devices can now be accessed and manipulated by applications that use standard UNIX `read(2)` and `write(2)` system calls, and without writing a special kernel driver.

Additional features include the following:

- Gives applications access to raw device data and device status
- Supports control, bulk, and interrupt (in and out) transfers

For more information, refer to the `ugen(7D)` man page and the USB Driver Development Kit at <http://developers.sun.com/>.

- **Digi Edgeport USB support** – Provides support for several Digi Edgeport USB-to-serial-port converter devices.
 - New devices are accessed as `/dev/term/[0-9]*` and `/dev/cua/[0-9]*`.
 - USB serial ports are usable as any other serial port would be, except that they cannot serve as a local serial console. The fact that their data is run through a USB port is transparent to the user.

For more information, see the `usbser_edge(7D)` man page, or go to <http://www.digi.com> and <http://www.sun.com/io>.

- **Documentation and binary support for user-written kernel and userland drivers** – A Solaris 9 USB Driver Development Kit (DDK) is available, and its documentation is applicable to the Solaris 10 OS. For up-to-date information about USB driver development, including information on the Driver Development Kit, go to <http://developers.sun.com/>.

EHCI and OHCI Drivers

This description is new in the Solaris Express 9/03 release. For Solaris 9 users, this feature is new in the Solaris 9 4/04 release.

Features of the EHCI driver include the following:

- Compliance with enhanced host controller interface that supports USB 2.0.
- Support for high-speed control, bulk, and interrupt transfers.
- Currently, no support is available for high-speed isochronous transactions. For example, you cannot connect USB 1.x devices to a 2.0 hub that is connected to a USB 2.0 port.

Note the following when using EHCI and OHCI controlled devices:

- A USB 2.0 host controller has one high-speed Enhanced Host Controller (EHCI) and one or more low-speed or full-speed OpenHCI Host Controller (OHCI) embedded controllers. Devices connected to a USB 2.0 port are dynamically assigned to either an EHCI or OHCI controller, depending on whether they support USB 2.0.
- If USB 2.0 and USB 1.x devices are on the system, the EHCI and OHCI drivers “hand-off” device control, depending on the type of device that is connected to the system.

Note – USB 2.0 storage devices connected to a port on a USB 2.0 PCI card, and that were used with a prior Solaris release in the same hardware configuration, can change device names after upgrading to this release. This change occurs because these devices are now seen as USB 2.0 devices and are taken over by the EHCI controller. The controller number, *w* in `/dev/[r]dsk/cwt.xdysz`, is changed for these devices.

For further information, see the man pages [ehci\(7D\)](#) and [ohci\(7D\)](#).

Logical Unit Reset

This feature is new in the Software Express pilot program. This feature is included in the Solaris 10 3/05 release.

The Solaris SCSI disk driver now supports SCSI logical unit reset. This feature improves error recovery processing on multi-LUN SCSI target devices. This feature can reset a logical unit without affecting other logical units on that device.

The ability to reset LUNs is particularly useful in conjunction with multiple pathing and with storage units that provide a large number of LUNs. For more information, see the [scsi_reset\(9F\)](#) and the [tran_reset\(9E\)](#) man pages.

Networking Enhancements

This section describes all networking enhancements in the Solaris 10 3/05 release that are new or have been enhanced since the Solaris 9 OS was originally distributed in May 2002.

Virtual IP Source Address Selection

This feature is new in the Solaris 10 3/05 release.

Virtual IP source address selection enables a system administrator to specify an IP source address to be used for packets that are routed through a particular network interface. This source address can be hosted on the just-introduced virtual network interface (`vni`), which is immune to hardware failures. Alternatively, the source address can be hosted on the loopback interface.

You can use virtual IP source address selection in conjunction with IP routing protocols to provide multipathing (that is, redundancy) at the network layer, beyond the first router. Currently, this form of multipathing works in conjunction with the RIPv2 routing protocol that is available in the `in.routed` daemon.

For more information about the virtual IP source address selection feature, see the [ifconfig\(1M\)](#) and [vni\(7d\)](#) man pages.

Stream Control Transmission Protocol

This feature is new in the Solaris Express 8/04 release.

Stream Control Transmission Protocol (SCTP) is a reliable transport protocol that is now included in the Solaris Operating System's TCP/IP protocol stack. SCTP provides services that are similar to TCP. However, SCTP supports connections between endpoints that are multihomed, that is, with more than one IP address. The support for multihoming makes SCTP a popular transport protocol for telephony applications. SCTP also supports multistreaming and partial reliability.

The SCTP protocol does not require additional configuration after the Solaris 10 OS is installed. However, you might need to add service definitions, so that particular applications can run over SCTP.

For information about configuring SCTP, refer to the *System Administration Guide: IP Services*.

Zebra Multiprotocol Routing Suite

This feature is new in the Solaris Express 8/04 release.

The Solaris 10 OS now includes the Zebra multiprotocol routing suite. This suite includes Open Source Zebra 0.92a routing software with bug fixes for Sun platforms. Now system administrators can use the well-known open-source routing protocols RIP, BGP, and OSPF for administering their Solaris based networks. Moreover, the OSPF daemon can be used for high network availability on multihomed servers. The Zebra packages contain these protocols and the `zebraadm` administration tool.

Refer to the `/etc/sfw/zebra/README.Solaris` file for configuration information and other details.

IPsec and NAT Traversal

This feature is new in the Solaris Express 8/04 release.

IKE can now initiate IPsec security associations from behind a Network Address Translation (NAT) box. Only the ESP protocol over an IPv4 network is allowed. Additionally, IPsec security associations for traffic that traverses a NAT cannot be accelerated with the Solaris Crypto Accelerator 4000 board. IKE acceleration is unaffected.

For more information, see the [ipseckey\(1M\)](#) man page.

Enhancement to the `nfsmapid` Daemon

This feature is new in the Solaris Express 8/04 release and updated in the Solaris 10 3/05 release.

In NFS version 4, the `nfsmapid` daemon provides a mapping from a numeric user identification (UID) or a numeric group identification (GID) to a string representation, as well as the reverse. The string representation is used by the NFS version 4 protocol to represent `owner` or `owner_group`.

For example, the UID 123456 for the user, `known_user`, that is operating on a client that is named `system.anydomain.com`, would be mapped to `known_user@anydomain.com`. The NFS client sends the string representation, `known_user@anydomain.com`, to the NFS server. The NFS server maps the string representation, `known_user@anydomain.com`, to the unique UID 123456. `nfsmapid` uses the `passwd` and `group` entries in the `/etc/nsswitch.conf` file to determine which database will be consulted to perform the mappings.

For `nfsmapid` to work properly, clients and servers on NFS version 4 must have the same domain. To ensure that clients and servers have the same domain, `nfsmapid` configures the domain by following these strict precedence rules:

- The daemon first checks the `/etc/default/nfs` file for a value that has been assigned to the `NFSMAPID_DOMAIN` keyword. If a value is found, the assigned value takes precedence over any other settings. The assigned value is appended to the outbound attribute strings and is compared against inbound attribute strings.
- If no value has been assigned to `NFSMAPID_DOMAIN`, then the daemon checks for a domain name from a DNS TXT record on a DNS name server. To find a specific DNS record, `nfsmapid` relies on the [`resolv.conf\(4\)`](#) configuration file.
- If no DNS TXT record provides a domain name, then by default the `nfsmapid` daemon uses the local DNS domain.

Note – The use of DNS TXT records is preferred. Configure the `_nfsv4idmapdomain` TXT record on DNS servers that provide domains for NFS version 4 clients and servers. TXT records provide better support for scaling issues and provide a single point of control.

If your network includes multiple DNS domains, but has only a single UID and GID namespace, all clients must use one value for `NFSMAPID_DOMAIN`. For sites that use DNS, `nfsmapid` resolves this issue by obtaining the domain name from the value that you assigned to `_nfsv4idmapdomain`. If your network is not configured to use DNS, during the first system boot the Solaris OS uses the [`sysidconfig\(1M\)`](#) utility to provide prompts for an NFS version 4 domain name.

For more information, see the `nfsmapid(1M)` and `sysidtool(1M)` man pages. See also the *System Administration Guide: Network Services*.

sendmail Version 8.13

Introduced in the Solaris Express 8/04 release, sendmail version 8.13 is the default in the Solaris 10 OS. Although this new version of sendmail provides many new features, the `FallBackSmartHost` option is the most significant addition.

Because of the `FallBackSmartHost` option you no longer need to use `main.cf` and `subsidiary.cf`. The `main.cf` file was used in environments that supported MX records. The `subsidiary.cf` file was used in environments without a fully operative DNS. In such environments a smart host was used instead of MX records.

The `FallBackSmartHost` option provides unified configuration. This option operates like an MX record of last possible preference for all environments. To ensure that mail gets delivered to clients, this option, if enabled, provides a well-connected, or “smart,” host that serves as a backup or failover for MX records that fail.

sendmail Version 8.13 also provides the following:

- Additional command-line options
- Additional and revised configuration file options
- Additional and revised FEATURE declarations

For more information, see the *System Administration Guide: Network Services*.

sendmail Version 8.12 Uses TCP Wrappers

This feature is new in the Solaris Express 9/03 release.

Note – In the Solaris Express 8/04 release, sendmail version 8.13 is the default. See “[sendmail Version 8.13](#)” on page 225.

TCP wrappers provide a way of implementing access controls by checking the address of a host that is requesting a particular network service against an access control list. Requests are granted or denied, accordingly. Besides providing this access control mechanism, TCP wrappers also log host requests for network services, which is a useful monitoring function. Examples of network services that might be placed under access control include `rlogind`, `telnetd`, and `ftpd`.

In this Solaris release, version 8.12 of sendmail now enables the use of TCP wrappers. This check does not bypass other security measures. By enabling TCP wrappers in sendmail, a check has been added to validate the source of a network request before the request is granted. See the `hosts_access(4)` man page.

Note – The Solaris 9 release added support for TCP wrappers in [inetd\(1M\)](#) and [sshd\(1M\)](#).

Sun Java System Message Queue

The Solaris Express 6/04 release introduced Sun Java System Message Queue 3.5 SP1 Platform Edition. This feature is included in the Solaris 10 3/05 release. This version replaces Sun Java System Message Queue (MQ) 3.0.1, which was previously introduced in the Software Express pilot program.

Note – For Solaris 9 users, the Message Queue 3.0.1 is new for the SPARC platform in the Solaris 9 12/02 release. In the Solaris 9 8/03 release, this feature was available for the x86 platform.

Sun Java System Message Queue 3.5 SP1 Platform Edition is an affordable, standards-based, high-performance Messaging System that integrates disparate IT systems. This Message Queue (MQ) is Java Messaging Services (JMS) 1.1 compliant and supports web services messaging through JAXM (SOAP 1.1 with Attachments).

The new 3.5 release, as compared to the previous 3.0.1 release, includes the following new features:

- **C Messaging API** – Native-to-the-wire C-API for connectivity to legacy C/C++ applications. TCP and SSL transports are supported.
- **Client Connection Failover** – On loss of connection, client automatically reconnects to a different broker in the cluster.
- **Advanced Remote Monitoring Capabilities** – JMS-based API enables monitoring of broker statistics, destination statistics, and VM statistics.
- **Support for Sun Cluster in Java Enterprise System** – Sun Cluster Agent for MQ, available in the Java Enterprise System, uses a file-based datastore with high availability (HA) to provide HA functionality to MQ.
- **J2EE 1.4 compatibility** – J2EE Connector Architecture 1.5 support enables MQ 3.5 to be connected to any J2EE 1.4 compliant Application Server through the MQ Resource Adapter.
- **Dynamic Message Flow Control** – Flow control enables management of throughput and load in the System. Control options include Reject Oldest, Reject Newest, Reject Low Priority, and Flow Control.
- **Local Destinations and Cluster Delivery Policies** – Policy configuration enables throughput optimization for message delivery to brokers in a cluster.

For new feature details, refer to the *Sun Java System Message Queue Release Notes* at <http://docs.sun.com>.

Sun Java System Application Server

Sun Java System Application Server Platform Edition 8 is new in the Solaris Express 6/04 release. This feature is included in the Solaris 10 3/05 release.

Previously, Sun Java System Application Server 7 is new in the Software Express pilot program for SPARC platforms, and in the Solaris Express 9/03 release for x86 platforms. Sun Java System Application Server Platform Edition 8 replaces Edition 7 in the Solaris Express 6/04 release.

Note – For Solaris 9 users, version 7 of the Application Server is new for SPARC platforms in the Solaris 9 12/02 release, and for x86 platforms in the Solaris 9 12/03 release.

Sun Java System Application Server Platform Edition 8 is the J2EE 1.4 compatible application container from Sun Microsystems that is designed for developers and departmental deployments.

This edition provides the following features:

- **J2EE 1.4 compatible** – Supports the latest J2EE technology standards and ensures application portability.
- **Integrated support for JavaServer Faces** – Enables developers to use the latest presentation layer technology, JavaServer™ Faces.
- **Administration GUI and Log Viewer** – Provides improved display and search mechanism for the Log Viewer, thus enhancing administrator productivity.
- **Deploytool** – A GUI-based tool supplements text editors, enabling developers to assemble J2EE components and deploy J2EE applications.
- **Improves Developer productivity** – Improvements made specifically to increase developer productivity include a reduced process count, a reduced memory footprint, an increased server startup speed, and an increased deployment speed.

For important information about this product, see the Sun Java System Application Server product information at <http://www.sun.com/software/>. For new feature details, refer to the *Sun Java System Application Server Platform Edition 8 Release Notes* at <http://docs.sun.com/>.

Using CacheFS With NFS Version 4

This feature is new in the Solaris Express 6/04 release.

The cache file system, CacheFS™, is a generic, nonvolatile caching mechanism. CacheFS improves the performance of certain file systems by utilizing a small, fast local disk. You can improve the performance of the NFS environment by using CacheFS.

CacheFS works differently with different versions of NFS. For example, if both the client and the back file system are running NFS version 2 or version 3, the files are cached in the front file

system for access by the client. However, if both the client and the server are running NFS version 4, the functionality is as follows. When the client makes the initial request to access a file from a CacheFS file system, the request bypasses the front (or cached) file system and goes directly to the back file system. With NFS version 4, files are no longer cached in a front file system. All file access is provided by the back file system. Also, because no files are being cached in the front file system, CacheFS-specific mount options, which are meant to affect the front file system, are ignored. CacheFS-specific mount options do not apply to the back file system.

Note – The first time you configure your system for NFS version 4, a warning appears on the console to indicate that caching is no longer performed.

For further information about NFS version 4, see *System Administration Guide: Network Services*.

Enhancement for `vacation` Utility

This feature is new in the Solaris Express 5/04 release.

The `vacation` utility has been enhanced to enable a user to specify which incoming messages receive autogenerated replies. With this enhancement, the user can avoid sharing confidential or contact information with unknown people. Messages from “spammers” or unknown people would not receive a reply.

This enhancement works by matching an incoming sender's email address to a list of domains or email addresses in a `.vacation.filter` file. This file is created by the user and is located in the user's home directory. If a domain or address match is found, a reply is sent. If no match is found, no reply is sent.

For more details, see the `vacation(1)` man page.

MILTER, `sendmail`'s New Mail Filter API

The feature is new in the Solaris Express 4/04 release.

MILTER, `sendmail`'s new mail filter API, permits third-party programs to access mail messages as they are being processed in order to filter meta-information and content. This functionality, introduced in the Solaris 10 3/05 release, requires the following:

- `sendmail` binary must be compiled with `-DMILTER`, which has been available since the Solaris 9 release.
- The file `/usr/lib/libmilter.so`, which is in the Solaris 10 OS, must be available.
- These files, `/usr/include/libmilter/mfapi.h` and `/usr/include/libmilter/mfdef.h`, must be available.

Both these files are included in the Solaris 10 OS.

Thus, with the Solaris 10 3/05 release, the user can build the filter and configure sendmail to use it.

For further information about sendmail, see *System Administration Guide: Network Services*.

IPv6 Advanced Sockets API

This feature is new in the Solaris Express 1/04 release.

The IPv6 Advanced Sockets API updates the Solaris Sockets API to meet the current version of RFC 2292. The advanced API provides the functionality needed to manipulate ICMP packets, obtain interface information, and manipulate IPv6 headers.

For further information, see the *Programming Interfaces Guide*.

Contents of /usr/lib/mail Have Moved to /etc/mail/cf

This feature is new in the Solaris Express 1/04 release.

The contents of the /usr/lib/mail directory, which might be in a read-only file system, are now in the /etc/mail/cf directory, which is writable. This change better supports m4 configuration. Note, however, these exceptions. The shell scripts /usr/lib/mail/sh/check-hostname and /usr/lib/mail/sh/check-permissions are now in the /usr/sbin directory. For backward compatibility, symbolic links point to each file's new location.

For further information, see the *System Administration Guide: Network Services*.

Added IPv6 Functionality During Solaris Installation

This feature is new in the Solaris Express 11/03 release.

Several new IPv6 functions are started when you select Enable IPv6 during Solaris installation.

- The /etc/nsswitch.conf file policies for the hosts database and ipnodes repositories are synchronized so that the same naming repositories are searched for ipnodes and hosts. Now, hosts can resolve any IPv6 addresses that might be in any of the ipnodes repositories.
- Destination address selection has been modified to avoid using an IPv6 address for a remote host if no IPv6 routes serve that host. Instead, an IPv4 address is used, to avoid any delays when connecting to remote hosts.

For example, consider the case of an IPv6-enabled host on a network with no IPv6 router. Without the presence of a router, a host has no knowledge of IPv6 routes beyond the local link. Previously, the host experienced timeouts when trying to connect to a preferred IPv6 address. With the new feature for destination address selection, the host prefers an IPv4 destination address. This feature eliminates time-out problems.

For further information, see the *System Administration Guide: IP Services*.

IPv6 Temporary Addresses

The feature is new in the Solaris Express 11/03 release.

By default, the interface ID of an IPv6 address is autoconfigured with the interface's hardware-specific address. For example, if the interface is an Ethernet card, its interface ID is autoconfigured from the interface's MAC address. However, some system administrators might need to keep private the hardware-specific addresses of one or more interfaces on a node.

The IPv6 Temporary Addresses feature implements the privacy extensions standard that is defined in RFC 3041, "Privacy Extensions for Stateless Autoconfiguration in IPv6." The temporary address feature enables administrators to assign randomly generated, modified EUI-64 format interface IDs to interfaces of an IPv6 node. In addition, the administrators can specify a time limit for the lifetime of a temporary address. After configuration, the IPv6 daemon `in.ndpd` automatically generates the temporary interface ID for the interface, in addition to the automatically generated, MAC address-based interface ID.

For further information, see the [System Administration Guide: IP Services](#).

`routeadm` Command

This feature is new in the Solaris Express 9/03 release.

The new `routeadm` command enables system administrators to configure IP forwarding and routing on all interfaces of a system. Any settings that are established through `routeadm` override system defaults that are read from configuration files at boot time.

The `routeadm` command provides options for enabling or disabling the global packet-forwarding function on all IPv4 or IPv6 interfaces of a system. `routeadm` can also be used to set up a system as a router, by turning on routing daemons for all system interfaces. In an IPv6 environment, `routeadm` can also be used to enable or disable dynamic routing for a host.

For further information, see the [`routeadm\(1M\)`](#) man page and the [System Administration Guide: IP Services](#).

TCP Multidata Transmit

Multidata transmit (MDT) is available only for systems that run a 64-bit kernel. This feature is new in the Software Express pilot program and in Solaris 9 8/03 release. This feature is included in the Solaris 10 3/05 release.

Multidata transmit enables the network stack to send more than one packet at one time to the network device driver during transmission. Use of this feature reduces the per-packet processing costs by improving the host CPU utilization or network throughput.

The MDT feature is only effective for device drivers that support this feature.

MDT is enabled by default. MDT can be disabled by including the following line in the `/etc/system` file:

```
# ndd -set /dev/ip ip_multidata_outbound 0
```

For further information, see the [ip\(7P\)](#) and the [nnd\(1M\)](#) man pages.

See also the [STREAMS Programming Guide](#).

`ifconfig` router Option

This feature is new in the Software Express pilot program. This feature is included in the Solaris 10 3/05 release.

The new router option of `ifconfig` allows you to configure IP packet forwarding on individual interfaces. `ifconfig router` and `ifconfig -router enable` or `disable` IP packet forwarding, respectively, for both IPv4 and IPv6 interfaces. The router option sets the `IFF_ROUTER` interface flag.

These new options replace the `nnd` variables `interface-name:ip_forwarding` and `interface-name:ip6_forwarding` for configuring IP packet forwarding on individual interfaces. Though now obsolete, the `nnd` variables remain in this Solaris release for backward compatibility. You can still use `ip_forwarding` and `ip6_forwarding` without the `interface-name` prefix to configure IP forwarding for all interfaces on a system.

For detailed information, refer to the [ifconfig\(1M\)](#), [ip\(7P\)](#), and [ip6\(7P\)](#) man pages.

Internet Protocol Version 6 (IPv6) Default Address Selection

This feature is new in the Software Express pilot program. This feature is included in the Solaris 10 3/05 release.

The Solaris 10 Operating System provides a documented and deterministic algorithm for doing IPv6 default source and destination address selection. This feature gives system administrators the limited ability to change address selection precedence through use of a policy table.

The IPv6 default address selection feature is a standardized method for IPv6 source and destination address selection.

The selection mechanisms can be configured by using a policy table. For example, you can edit the policy table to give higher precedence to a particular address prefix. Thereafter, addresses that are within the prefix are sorted before other addresses by name look-up APIs. You can also assign labels to source and destination prefixes in the policy table. This assignment makes sure that particular source addresses are only used with particular destination addresses.

To implement IPv6 default address selection, the Solaris Operating System now includes the `/etc/inet/ipaddrsel.conf` file and the `/usr/sbin/ipaddrsel` command. You use `ipaddrsel.conf` to edit the IPv6 default address policy table. Then you use `ipaddrsel` to commit the changes to the policy table.

Additionally, the `ifconfig` command now includes the “preferred” option. This option lets you designate a particular address to be used as the source address for all IPv6 communication.

For detailed information, refer to the `ipaddrsel.conf(4)`, `ipaddrsel(1M)`, and the `ifconfig(1M)` man pages.

For further information, see the *System Administration Guide: IP Services*.

Disable NFS and automount Services

This feature is new in the Software Express pilot program. This feature is included in the Solaris 10 3/05 release.

The following NFS daemons are typically started by default at boot time by the `rc` scripts: `nfsd`, `mountd`, `statd`, `lockd`, and the automount daemon, `automountd`. Now, if a machine does not require NFS and automount services, the scripts do not start the NFS daemons and the automount daemon.

The following describes the new behavior:

- The automount utility, which is called by `/etc/init.d/autofs`, now starts `automountd` at boot time only if the automount maps have a valid entry.
- `/etc/init.d/nfs.server` starts `mountd`, `nfsd`, `statd`, and `lockd` at boot time only if the machine has any NFS exports.
- `/etc/init.d/nfs.client` starts `statd` and `lockd` only if `/etc/vfstab` includes NFS file systems.

NFS and automount services might not be started at boot time. The following commands can start these services when required by a machine:

- The automount utility can start the `automountd` daemon.
- The `mount` command with the `-F nfs` option can start the `lockd` and `statd` daemons. The `automountd` daemon also can start `lockd` and `statd`.
- The `share` command with the `-F nfs` option can start the `nfsd`, `mountd`, `lockd`, and `statd` daemons.

This mechanism for disabling NFS and automount services provides the following benefits:

- Extra security that comes from not running unnecessary daemons on a machine.
- A simplified process for exporting file systems. The `nfsd` and `mountd` daemons, and, if necessary, `lockd` and `statd`, are started by the `share` command with the `-nfs` option. So, you no longer have to edit the `/etc/dfs/dfstab` file and then invoke the `/etc/init.d/nfs.server` program. This new behavior permits an NFS export to be configured with a single command, without editing any configuration files. However, if the system reboots, such exports are not resumed automatically, unless the exports are included in the `/etc/dfs/dfstab` file.

For more information, refer to the man pages for [mountd\(1M\)](#), [lockd\(1M\)](#), [statd\(1M\)](#), and [nfsd\(1M\)](#).

For further information, see also the *System Administration Guide: Network Services*.

Internet Protocol Version 6 (IPv6) 6to4 Router

This feature is new in the Software Express pilot program and in the Solaris 9 4/03 release. This feature is included in the Solaris 10 3/05 release.

IPv6 networks can now transfer packets over Internet Protocol Version 4 (IPv4) networks by configuring one or more routers to support a 6to4 tunnel. System administrators can use 6to4 tunnels as a transitional method for migrating their networks from IPv4 to IPv6. This feature implements RFCs 3056 and 3068.

For further information on IPv6, see the *System Administration Guide: IP Services*.

Packet Tunneling Over IPv6

This feature is new in the Software Express pilot program and in the Solaris 9 9/02 release. This feature is included in the Solaris 10 3/05 release.

This feature enables tunneling over IPv6 for both IPv4 over IPv6 tunnels and IPv6 over IPv6 tunnels. IPv4 packets or IPv6 packets can be encapsulated in IPv6 packets.

For further information, see the *System Administration Guide: IP Services*.

Hosting Multiple Web Sites on a Single Solaris Machine

This feature is new in the Software Express pilot program and in the Solaris 9 12/02 release. This feature is included in the Solaris 10 3/05 release.

The Solaris Network Cache and Accelerator (NCA) kernel module now supports multiple instances of a web server. This support enables you to use a Solaris machine to perform Internet protocol (IP) address-based virtual web hosting. The Solaris software uses a single configuration file, `/etc/nca/ncaport.conf`, to map NCA sockets to IP addresses.

For further information, see the [ncaport.conf\(4\)](#) man page.

IP Quality of Service

This feature is new in the Software Express pilot program and in the Solaris 9 9/02 release. This feature is included in the Solaris 10 3/05 release.

IP Quality of Service (IPQoS) is a new feature in the Solaris Operating System. IPQoS enables system administrators to provide different levels of network service to customers and to critical applications. By using IPQoS, the administrator can set up service-level agreements. These agreements provide an Internet service provider's (ISP) clients with varying levels of service that are based on a price structure. A company could also use IPQoS to prioritize among applications so that critical applications get a higher quality of service than less critical applications.

For further information, see the *System Administration Guide: IP Services*.

User Selector for Internet Protocol Quality of Service (IPQoS)

This feature is new in the Software Express pilot program and in the Solaris 9 8/03 release. This feature is included in the Solaris 10 3/05 release.

The Solaris IPQoS feature now includes the user selector, which supplements the existing `uid` selector. The user selector enables you to specify a user name or `userID` as criteria in a filter clause in the `ipqos.conf` file. Previously, the `uid` selector only accepted a `userID` as a value. The following filter clause from an `ipqos.conf` file shows the user selector:

```
filter {
    name myhost;
    user root;
}
```

For information about filters and selectors, refer to the [ipqosconf\(1M\)](#) man page.

See also the *System Administration Guide: IP Services*.

Routing Information Protocol version 2 (RIPv2)

This feature is new in the Software Express pilot program and in the Solaris 9 9/02 release. This feature is included in the Solaris 10 3/05 release.

Solaris system software now supports Routing Information Protocol version 2 (RIPv2).

RIPv2 adds Classless Inter-Domain Routing (CIDR) and Variable-Length Subnet Mask (VLSM) extensions to the RIPv1 protocol. Message Digest 5 (MD5) extensions protect routers

against intentional misdirection by malicious users. The new `in.routed` implementation also includes a built-in Internet Control Message Protocol (ICMP) Router Discovery (RFC 1256) mechanism.

RIPv2 supports multicast if the point-to-point links are enabled with multicast. RIPv2 also supports unicast. If you configure a broadcast address by using the `/etc/gateways` file, then RIPv2 supports broadcast.

For information on how to configure RIPv2, see the [`in.rdisc\(1M\)`](#), [`in.routed\(1M\)`](#), and [`gateways\(4\)`](#) man pages.

Developer Tool Enhancements

This section describes all development tools in the Solaris 10 3/05 release that are new or have been enhanced since the Solaris 9 OS was originally distributed in May 2002. DTrace dynamic tracing facility is a feature of particular importance.

Developers should also review these new features in the security and system administration sections:

- “Kernel Modular Debugger” on page 169
- “OpenSSL and OpenSSL PKCS#11 Engine” on page 199

Dynamic Tracing Facility

This feature is new in the Solaris Express 11/03 release.

The Solaris DTrace is a comprehensive dynamic tracing facility that gives Solaris users, administrators, and developers a new level of observability into the kernel and user processes. For further information, see “DTrace Dynamic Tracing Facility” on page 162.

GCC Version 3.4.3

This feature is new in the Solaris 10 3/05 release.

GCC is the open source “C” compiler from the GNU Software Foundation. Build tools include `gmake`, `bison`, `binutils`, `gnuM4`, and `flex`.

Perl Version 5.8.4

This feature is new in the Solaris Express 8/04 release.

A new, default version of the Practical Extraction and Report Language (Perl) is available in the Solaris 10 OS. The new default version of Perl is version 5.8.4. You need to reinstall any modules that you have manually installed after upgrading to the Solaris 10 software. For more information about installing modules, see “Perl 5” in the *System Administration Guide: Network Services*.

For more information about Perl, see the [perl\(1\)](#) man page.

Per-thread Mode Enhancement

This feature is new in the Solaris Express 8/04 release.

The `ttruss`, `pstack`, and `pflags` commands have been updated to enable the user to specify individual threads within a process or core. Developers and system administrators can better understand the behavior of large multithreaded applications, and target specific threads for debugging.

For further information, see the following man pages:

- [ttruss\(1\)](#)
- [pstack\(1\)](#)
- [pflags\(1\)](#)

USB End-User Device Support Enhancements

These enhancements are new in the Solaris Express 5/04 release. New programming interfaces were added in the Solaris Express 8/04 release.

All current Edgeport USB-to-serial devices now work with this Solaris release. USB 1.1 audio and other isochronous devices now work behind USB 2.0 high-speed hubs.

New Programming Interfaces

These interfaces are new in the Solaris Express 8/04 release.

Universal Serial Bus Architecture (USB) 2.0 programming interfaces are being released publicly as part of the Solaris 10 OS. These interfaces are documented as `usb_*` functions and structures in man page sections 9F and 9S. See also *Writing Device Drivers*.

The USB framework now is called USB 2.0. USB drivers that were written to USB 1.0 interfaces in prior DDks are only binary compatible with the Solaris 10 software. These drivers are not source compatible with the Solaris 10 software.

ls Enhancements

This feature is new in the Solaris Express 8/04 release.

The `ls` command can now display file times to the second or to the nanosecond. See [“Enhancements to ls Command” on page 170](#).

New Functions for Converting Strings

This feature is new in the Solaris Express 7/04 release.

Two new functions for converting strings to long integers are now available to kernel module developers. The `ddi_strtol()` and `ddi_strtoul()` functions enable conversion from strings to long integers and unsigned long integers respectively. The new functions provide flexible input, DDI compliancy, and increased error detection.

For more information, refer to the [ddi_strtol\(9F\)](#) and [ddi_strtoul\(9F\)](#) man pages.

Java Support for `pstack` Command

This feature is new in the Solaris Express 7/04 release.

The `pstack` command has been enhanced to display Java frames when using the latest version of Java. For each Java frame in the call stack, the function and line number displays, if available.

See the [pstack\(1\)](#) man page.

New Mechanisms for the Solaris Cryptographic Framework

This feature is new in the Solaris Express 6/04 release.

The Solaris Cryptographic Framework now supports the following mechanisms for the Secure Sockets Layer (SSL) and Transport Layer Security (TLS) protocols:

- CKM_SSL3_PRE_MASTER_KEY_GEN
- CKM_SSL3_MASTER_KEY_DERIVE
- CKM_SSL3_KEY_AND_MAC_DERIVE
- CKM_SSL3_MASTER_KEY_DERIVE_DH
- CKM_TLS_PRE_MASTER_KEY_GEN
- CKM_TLS_MASTER_KEY_DERIVE
- CKM_TLS_KEY_AND_MAC_DERIVE
- CKM_TLS_MASTER_KEY_DERIVE_DH

The Solaris Cryptographic Framework is an architecture that enables applications in the Solaris system to use or provide cryptographic services. All interactions with the framework are based on the PKCS#11 Cryptographic Token Interface (Cryptoki) from RSA Laboratories, the research arm of RSA Security, Inc.

For more information, see “PKCS #11 Functions: `C_GetMechanismList`” in the *Solaris Security for Developers Guide*.

Retail and Nonretail Options for Providers in Solaris Cryptographic Framework

This feature is new in the Solaris Express 6/04 release.

Vendors of cryptographic providers that connect to the Solaris Cryptographic Framework have more flexibility in requesting certificates from Sun Microsystems. Certificates now support both retail and nonretail export distribution.

Retail encryption products are those products that the U.S. government permits to be shipped to all countries. Retail products cannot be shipped, however, to designated nations that the U.S. government considers to be security threats. Nonretail encryption products are those products that the U.S. government has permitted for domestic use only or for countries that the government has exempted.

For more information, see the [elfsign\(1\)](#) man page and Appendix F, “Packaging and Signing Cryptographic Providers,” in the *Solaris Security for Developers Guide*.

Linkers and Libraries Updates

This description is new in the Software Express pilot program and revised in the Solaris Express 5/04 release. Linker and libraries updates were also introduced in the Solaris 9 12/02 and subsequent releases. These updates are included in the Solaris 10 3/05 release.

The Solaris 10 OS includes linker-editor features such as string table compression, unreferenced section elimination, and unreferenced dependency detection. For complete information on the most recent feature enhancements, see Appendix D, “Linker and Libraries Updates and New Features,” in the *Linker and Libraries Guide*.

Enhancements introduced in the Solaris Express 5/04 release include the following:

- A restructuring of the file system has moved many components from under `/usr/lib` to `/lib`. The default search paths for both the link-editor and runtime linkers have been changed accordingly.
- System archive libraries are no longer provided. Therefore, the creation of a entirely statically-linked executable is no longer possible.
- Greater flexibility for defining alternative dependencies is provided with the `-A` option for the `crle` command.

Enhancements introduced in the previous Software Express pilot program include the following:

- Greater flexibility in defining the hardware and software requirements of ELF objects is provided with the link-editors.
- The runtime link-auditing interface `la_objfilter()` has been added.
- Shared-object filtering has been extended to provide filtering on a per-symbol basis.
- Thread-local storage is provided.
- The `-z ignore` option has been extended to eliminate unreferenced sections during a link-edit. See the [ld\(1\)](#) man page.
- Greater flexibility in defining a symbol's visibility is provided with the “protected” mapfile directive.

- `dlopen(3DL)` and `dlsym(3DL)` look-up semantics have been expanded with a new mode, `RTLD_FIRST`.
- Unreferenced dependencies can be determined by using the `ldd` utility. See the `-U` option in the `ldd(1)` man page.

Layered Driver Interfaces

This feature is new in the Solaris Express 3/04 release.

The Solaris 10 Operating System has been enhanced to enable kernel modules to perform device access operations such as open, read, and write a device. This system also enables you to determine which devices are provided through a new set of publicly available Layered Driver Interfaces, referred to as the “LDI.”

Driver developers can use the LDI interfaces to access character, block, or STREAMS devices directly from within the Solaris kernel. Application developers can use the LDI interfaces to display device layering information. This new architecture also provides administrators with observability into device usage inside the kernel. For further information, see the `ldi_*(9F)` and the `di_*(3DEVINFO)` man pages.

The `prtconf` and `fuser` utilities have been enhanced to include the following capabilities:

- “Device Layering” through the `prtconf` command – This command displays device minor node and device usage information. The utility also shows which minor nodes a kernel module currently has open.
See the `prtconf(1M)` man page.
- “Device Usage” through the `fuser` command – This command displays information about users of a device. The command also shows what generic Solaris kernel subsystem or user process opens and accesses a device inside the Solaris kernel.
See the `fuser(1M)` man page.

The Layered Driver Interfaces begin with the prefix `ldi_`. These interfaces are used for device access and to obtain device information at the kernel level. Man pages are provided for the interfaces in section 9F. At the user level, a set of device information library interfaces provides for retrieving kernel device usage information within applications. Man pages are provided for the LDI `libdevinfo` interfaces within section 3DEVINFO. In addition, the `prtconf(1M)` and `fuser(1M)` man pages include information about displaying the kernel device usage information that is provided by the LDI architecture.

For further information, see Chapter 13, “Layered Driver Interface (LDI),” in the *Writing Device Drivers*.

Changes to `makecontext()` Function

This feature is new in the Solaris Express 3/04 release and in the Solaris 9 9/04 release.

The semantics of the `uc_stack` member of the `ucontext_t` structure have changes as they apply to inputs to the `makecontext(3C)` libc library function. Binary compatibility is preserved between previous versions of Solaris and the Solaris 10 OS.

Applications that use this interface must be updated before they are recompiled for the Solaris 10 OS. See the `makecontext(3C)` man page for further information.

Single UNIX Specification, Version 3

This feature is new in the Solaris Express 2/04 release.

This Solaris release conforms to the Single UNIX Specification, Version 3 (SUSv3). SUSv3 provides updates to POSIX.1-1990, POSIX.1b-1993, POSIX.1c-1996, POSIX.2-1992, and POSIX.2a-1992.

See “Single UNIX Specification, Version 3 Introduces Changes” in the *Solaris 10 Release Notes* for a detailed description of the impact of the SUSv3 updates for Solaris users.

Advanced API

This feature is new in the Solaris Express 1/04 release.

The IPv6 Advanced Sockets API updates the Solaris Sockets API to meet the current version of RFC 2292. See “IPv6 Advanced Sockets API” on page 229.

Simple Authentication and Security Layer for Developers

This feature is new in the Solaris Express 12/03 release.

Simple Authentication and Security Layer (SASL) provides developers of applications and shared libraries with interfaces for adding authentication, data integrity checking, and encryption to connection-based protocols.

SASL consists of the following items:

- Library, `libsasl`, which provides an API for applications that need authentication, privacy, and integrity services
- Service provider interface (SPI) for third-party plug-ins to add new authentication methods, name canonicalization rules, and property stores
- Header files for development
- Plug-ins that are supplied by Sun for these mechanisms:
 - EXTERNAL
 - PLAIN
 - CRAM-MD5
 - DIGEST-MD5
 - GSS-API

- GSS-SPNEGO

SASL enables the developer to write to a generic API without having to be concerned about the details of security mechanisms. When developed to use SASL appropriately, servers and clients can use new security mechanisms, naming and user canonicalization plug-ins, and auxprop plug-ins without recompilation.

SASL is described in RFC 2222. SASL is particularly appropriate for applications that use the following protocols that support SASL:

- IMAP
- SMTP
- ACAP
- LDAP

For more information about SASL, see the [libsasl\(3LIB\)](#) man page. See also the *Solaris Security for Developers Guide*.

Event Ports

This feature is new in the Solaris Express 12/03 release.

Event Ports is a framework that allows applications to generate and collect events from disjoint sources. The framework can retrieve events from multiple objects simultaneously without degrading overall performance.

For further information, see the [port_create\(3C\)](#) and [signal.h\(3HEAD\)](#) man pages.

Core File Content

The Solaris Express 12/03 release introduced enhancements to the `coreadm`, `gcore`, and `mdb` utilities. See “[Core File Content Enhancements](#)” on page 171.

Atomic Operations

This feature is new in the Solaris Express 10/03 release and has been enhanced in the Solaris 10 1/06 release.

Atomic operations provide APIs in `libc` that perform simple atomic operations quickly. This new feature permits applications to atomically update memory without using other synchronization primitives or platform-specific assembly language. The available operations include addition, the “and” Boolean function, and the “or” Boolean function.

For more information, see the [atomic_ops\(3C\)](#) man page.

Solaris WBEM File Changes

This feature description was updated in the Solaris Express 9/03 release.

There have been several changes to the Managed Object Format (MOF) files in the `/usr/sadm/mof` directory.

- The `Solaris_VM1.0.mof` file was revised to `Solaris_VM2.0.mof` and subsequently to `Solaris_VM3.0.mof`.
- The local file system classes were relocated from `Solaris_VM2.0.mof` to the new file, `Solaris_FS1.0.mof`. `Solaris_FS1.0.mof` defines classes that pertain to storage devices.
- Two of the providers, `Solaris_DiskDrive` and `Solaris_DiskPartition`, that were in the `Solaris_VM1.0.mof` file, were moved into the new `Solaris_DMGT.1.0.mof` file. The `Solaris_DMGT.1.0.mof` file contains classes that represent disks, disk partitions, and other device management classes.
- This release includes another new MOF file, `Solaris_NFS1.0.mof`. The `Solaris_NFS1.0.mof` file defines classes that pertain to NFS devices. This file contains the NFS classes from `Solaris_VM2.0.mof` as well as new classes for configuring and monitoring NFS shares (or “exports”) and mounts.

Privileges for Software Developers

This feature is new in the Software Express pilot program. This feature is included in the Solaris 10 3/05 release.

Processes need no longer run as root to have superuser capabilities. Instead, superuser capabilities can be parceled out by system administrators as discrete process rights. These process rights are implemented through privileges. Privileges enable developers to limit access to restricted operations and limit the periods for which privileges are in effect. The use of privileges can reduce the harm that formerly resulted if a privileged program was compromised. For compatibility, unmodified programs that run as root continue to have all privileges.

For general information about privileges, see “[Process Rights Management](#)” on page 195. For information about setting and getting privileges, see the `setppriv(2)` and `getppriv(2)` man pages. To learn more about manipulating privileges, see the `priv_str_to_set(3C)` and the `priv_addset(3C)` man pages.

For further information, see the *Solaris Security for Developers Guide*.

Solaris Cryptographic Framework for Developers

This feature is new in the Software Express pilot program. This feature is included in the Solaris 10 3/05 release.

The Solaris Cryptographic Framework provides cryptographic services to applications. Applications can access the framework through `libpkcs11(3LIB)` and at higher levels.

The Solaris Cryptographic Framework provides the following features for developers of applications that use encryption:

- User-level programming interfaces for various cryptographic functions. These interfaces cover, for example, encryption, decryption, message digests, and signing. The industry standard, RSA Security Inc. PKCS #11 Cryptographic Token Interface (Cryptoki), serves as the API.

The framework supports the following encryption algorithms:

- AES
 - DES/3DES
 - RC4
 - MD5
 - SHA-1
 - DSA
 - RSA
 - D-H
- User-level pluggable interfaces for Sun and third-party developers. These interfaces enable administrators to add new plug-ins from providers of encryption algorithms at the user level. Administrators can replace an existing provider with a different implementation. The user service provider interface (SPI) also uses the PKCS#11 standard. Tools for signing, packaging, and installing third-party binaries are provided.
 - An optimized software implementation of the most commonly used encryption and digital signing algorithms, such as AES, DES/3DES, and RSA. These implementations have been optimized for the SPARC platform and UltraSPARC platform.
 - An administrative CLI tool, `cryptoadm`, for adding or removing encryption plug-ins, setting cryptographic security policy, and other related administrative functions. See the [`cryptoadm\(1M\)`](#) man page.

See the following man pages: [libpkcs11\(3LIB\)](#), [pkcs11_softtoken\(5\)](#) and [pkcs11_kernel\(5\)](#). See also “Solaris Cryptographic Framework for System Administrators” on page 205.

Vendors of software or hardware cryptographic accelerators who are interested in supplying plug-ins to the Solaris cryptographic framework should contact Sun Microsystems for more details.

SPARC: 64-bit Package Changes for Software Developers

This feature is new in the Software Express pilot program. This feature is included in the Solaris 10 3/05 release.

In the Solaris 10 Operating System, packaging has been simplified with most 32-bit components and 64-bit components being delivered in a single package. See “[SPARC: 64-bit Package Changes](#)” on page 189.

SPNEGO Pseudo-Mechanism for GSS-API Applications

This feature is new in the Software Express pilot program. This feature is included in the Solaris 10 3/05 release.

This Solaris 10 OS includes a new GSS-API “pseudo-mechanism” for negotiating GSS-API security that is based on the SPNEGO protocol (IETF RFC 2478). Simple and Protected GSS-API Negotiation (SPNEGO) is most useful for applications that are based on GSS-API implementations which support multiple security mechanisms. SPNEGO can be applied when two applications use GSS-API to exchange data and do not know which mechanisms are supported by the other application.

SPNEGO is a pseudo-security mechanism that is represented by the following object identifier:

```
iso.org.dod.internet.security.mechanism.snego (1.3.6.1.5.5.2)
```

SPNEGO enables GSS-API peers to determine in-band whether their credentials share common GSS-API security mechanisms. If the mechanisms are shared, then the peers can select a common mechanism to establish the security context.

For further information, see the [mech\(4\)](#) and the [mech_spnego\(5\)](#) man pages. See also the *Solaris Security for Developers Guide*.

Locality Groups

This feature is new in the Software Express pilot program and in the Solaris 9 12/03 release. This feature is included in the Solaris 10 3/05 release.

The *Programming Interfaces Guide* now includes a chapter that explains the interfaces that interact with locality groups (lgroups). These interfaces can be used to help an application efficiently allocate CPU and memory resources. This capability results in improved performance on some systems.

Thread Stacks in pmap

This feature is new in the Software Express pilot program. This feature is included in the Solaris 10 3/05 release.

The `pmap` utility now labels threads' stacks so the stacks can be identified more easily.

See the [pmap\(1\)](#) man page for further information.

New DOOR_REFUSE_DESC Flag

This feature is new in the Software Express pilot program. This feature is included in the Solaris 10 3/05 release.

A new flag, `DOOR_REFUSE_DESC`, was added to the `door_create()` function. This new flag simplifies writing door servers which do not accept argument descriptors.

For further information, see the [door_create\(3DOOR\)](#) man page.

Stack Check APIs

This feature is new in the Software Express pilot program and in the Solaris 9 4/03 release. This feature is included in the Solaris 10 3/05 release.

Stack Check APIs allow for advanced interaction with stack-checking compiler support available in the Sun ONE Studio. These APIs should be used in applications that are compiled with stack checking enabled, and either manage their own stacks or attempt to detect their own stack overflows.

Developers who maintain their own thread library need to use the `setustack` interface to enable consumers of their library to compile with stack checking enabled.

See the [stack_getbounds\(3C\)](#), [stack_setbounds\(3C\)](#), and the [stack_inbounds\(3C\)](#) man pages.

Enhanced `crypt()` Function for Software Developers

This feature is new in the Software Express pilot program and in the Solaris 9 12/02 release. This feature is included in the Solaris 10 3/05 release.

The Software Express releases include new extensions to the `crypt()` function and introduce the `crypt_gensalt()` function. These enhancements allow administrators to change the algorithm that is used to obscure users' UNIX login passwords.

Modules are included for MD5 and Blowfish. The MD5 modules are at `crypt_sunmd5` and `crypt_bsdmd5`. The Blowfish module is at `crypt_bsd bf`.

Developers can create new modules for alternate password-obscuring algorithms. Application developers must use the `crypt_gensalt()` function instead of manually generating the salt string for passing to the `crypt()` function.

Modules for alternate algorithms are specified in the [crypt.conf\(4\)](#) file. The `module_path` field specifies the path to the shared library object that implements the two required functions:

- `crypt_gensalt_impl()` – Generates the salt string
- `crypt_genhash_impl()` – Generates the encrypted password

For further information, see the [crypt\(3C\)](#) and the [policy.conf\(4\)](#) man pages.

New Flags for `madvise()` Function

This feature is new in the Software Express pilot program and in the Solaris 9 12/02 release. This feature is included in the Solaris 10 3/05 release.

The `madvise()` function enables the kernel to optimize access to a user-defined region of memory. This Solaris release includes three new flags for the `madvise()` function:

- `MADV_ACCESS_LWP` – Gives a specified lightweight process (LWP) resource allocation priority
- `MADV_ACCESS_MANY` – Specifies an address range that is intensively used by processes across the machine
- `MADV_ACCESS_DEFAULT` – Resets an address range's access pattern to the system default

For further information on the `madvise()` function, see the [`madvise\(3C\)`](#) man page.

Memory Allocation With `libumem`

This feature is new in the Software Express pilot program and in the Solaris 9 4/03 release. This feature is included in the Solaris 10 3/05 release.

`libumem` is a user-mode (nonkernel mode) memory allocator library. `libumem` has features that enable you to debug memory leaks and other aberrations that involve memory usage.

This feature is used in the same way that a standard application binary interface (ABI) allocator, such as `malloc()`, is used. A user-mode application requests an arbitrary number of bytes of memory. Then a pointer is returned that is loaded with the address of the allocated memory.

For further information, see the [`libumem\(3LIB\)`](#) man page.

Smart Card Terminal Interfaces

This feature is new in the Software Express pilot program and in the Solaris 9 8/03 release. This feature is included in the Solaris 10 3/05 release.

Solaris smart card interfaces are a set of public interfaces for Smart Card Terminals. Card-terminal vendors can implement these interfaces in a user-level shared library to provide device-level support for their smart card terminals in Solaris. The set of Solaris smart card terminal interfaces is based on card terminal interfaces that are available as part of the Linux Smartcard framework. Card terminal support libraries from Linux can be ported to the Solaris environment with minimum effort. For further information on smart cards, see the [Solaris Smartcard Administration Guide](#).

Smartcard Middleware APIs

This feature is new in the Software Express pilot program and in the Solaris 9 9/02 release. This feature is included in the Solaris 10 3/05 release.

The Solaris Smartcard framework now provides low-level middleware APIs. These APIs can be used to exchange data with a smart card by using a smart card reader. The APIs can be used in platforms such as the Sun Blade™ and Sun Ray™ systems. Applications that are written in the Java language, or in C, can use these interfaces.

For more information, see the [libsmartcard\(3LIB\)](#) man page and the JavaDocs in `/usr/share/javadoc/smartcard`. See also the *Solaris Smartcard Administration Guide*.

File System Enhancements

This section describes all file system features in the Solaris 10 3/05 release that are new or have been enhanced since the Solaris 9 OS was originally distributed in May 2002.

NFS Version 4

This feature is new in the Software Express pilot program. In the Solaris Express 8/04 release, NFS version 4 became the default. This feature is included in the Solaris 10 3/05 release.

The Solaris 10 OS includes the Sun implementation of the NFS version 4 distributed file access protocol. This version is the next logical step in the evolution of NFS. The NFS version 4 protocol, specified in RFC 3530, was created under the auspices of the Internet Engineering Task Force (IETF). This version is designed to be both vendor neutral and operating system neutral.

NFS version 4 integrates file access, file locking, and mount protocols into a single, unified protocol to ease traversal through a firewall and improve security. The Solaris implementation of NFS version 4 is fully integrated with Kerberos V5, also known as SEAM, thus providing authentication, integrity, and privacy. NFS version 4 also enables the negotiation of security flavors to be used between the client and the server. With NFS version 4, a server can offer different security flavors for different file systems.

The Solaris implementation of NFS version 4 includes delegation, a technique by which the server can delegate the management of a file to a client. This technique can reduce the number of round-trip operations because the client is guaranteed that no modification can occur without the server informing the client. The protocol also includes operation compounding, which allows multiple operations to be combined into a single “over-the-wire” request.

For more information about NFS version 4, refer to Chapter 6, “Accessing Network File Systems (Reference),” in the *System Administration Guide: Network Services*.

UFS Logging Enabled by Default

This feature is new in the Solaris Express 4/04 release and in the Solaris 9 9/04 release.

Logging is now enabled by default for all UFS file systems except under the following conditions:

- When logging is explicitly disabled
- If insufficient file system space exists for the log

In previous Solaris releases, you had to enable UFS logging manually.

UFS logging packages into one transaction the multiple metadata changes that compose a complete UFS operation. Sets of transactions are recorded in an on-disk log, and then applied to the actual UFS file system's metadata.

UFS logging provides two advantages:

- If the file system is already consistent because of the transaction log, you might not have to run the `fsck` command after a system crash or an unclean shutdown.
- Starting in the Solaris 9 12/02 release, the performance of UFS logging improves or exceeds the level of performance of nonlogging file systems. This improvement can occur because a file system with logging enabled converts multiple updates to the same data into single updates. This capability reduces the number of overhead disk operations that are required.

For more information, see “What’s New in File Systems in the Solaris 10 Release?” in the *System Administration Guide: Devices and File Systems*. See also the `mount_ufs(1M)` man page.

Enhancements to the NFS Client

This feature is new in the Software Express pilot program and in the Solaris 9 12/03 release. This feature is included in the Solaris 10 3/05 release.

The following enhancements have improved the performance of the NFS client:

- Restrictions on wire transfer sizes have been relaxed. Now, the transfer size is based on the capabilities of the underlying transport. For example, the NFS transfer limit for UDP is still 32 Kbytes. However, because TCP is a streaming protocol without the datagram limits of UDP, maximum transfer sizes over TCP have been increased to 1 Mbyte.
- Previously, all write requests were serialized by both the NFS client and the NFS server. The NFS client has been modified to permit an application to issue concurrent writes, as well as concurrent reads and writes, to a single file. You can enable this functionality on the client by using the `forcedirectio` mount option. When you use this option, you are enabling this functionality for all files within the mounted file system. You could also enable this functionality on a single file on the client by using the `directio()` interface. Note that unless this new functionality has been enabled, writes to files are serialized. Also, if concurrent writes or concurrent reads and writes are occurring, then POSIX semantics are no longer being supported for that file.

- The NFS client no longer uses an excessive number of UDP ports. Previously, NFS transfers over UDP used a separate UDP port for each outstanding request. Now, by default, the NFS client uses only one UDP reserved port. However, this support is configurable. If the use of more simultaneous ports would increase system performance through increased scalability, then the system can be configured to use more ports. This capability also mirrors the NFS over TCP support, which has had this kind of configurability since its inception.

For further information, see the *System Administration Guide: Network Services*.

Multiterabyte UFS File Systems

Multiterabyte UFS file system support is available only for systems that run a 64-bit kernel. This feature is new in the Software Express pilot program and in the Solaris 9 8/03 release. This feature is included in the Solaris 10 3/05 release.

The Solaris 10 OS provides support for multiterabyte UFS file systems on systems that run a 64-bit Solaris kernel. Previously, UFS file systems were limited to approximately 1 terabyte (Tbyte) on both 64-bit systems and 32-bit systems. All UFS file system commands and utilities have been updated to support multiterabyte UFS file systems.

You can initially create a UFS file system that is less than one Tbyte. You can specify that the file system can eventually be grown to a multiterabyte file system by using the `newfs -T` command. This command sets the inode and fragment density to scale appropriately for a multiterabyte file system.

Support for a multiterabyte UFS file system assumes the availability of multiterabyte LUNs. These LUNs are provided as Solaris Volume Manager volumes, or as physical disks that are greater than one Tbyte.

Features of multiterabyte UFS file systems include the following:

- You can create a UFS file system to a maximum of 16 Tbytes in size.
- You can create a file system that is less than 16 Tbytes, which can later be increased in size to a maximum of 16 Tbytes.
- Multiterabyte file systems can be created on physical disks and on Solaris Volume Manager's logical volumes.
- UFS logging is enabled by default on file systems greater than 1 Tbyte. Multiterabyte file systems benefit from the performance improvements of having UFS logging enabled. Multiterabyte file systems also benefit from the availability of logging because the `fsck` command might not have to be run when logging is enabled.

Limitations of multiterabyte UFS file systems include the following:

- You cannot mount a file system that is greater than 1 Tbyte on a system that runs a 32-bit Solaris kernel.
- You cannot boot from a file system that is greater than 1 Tbyte on a system that runs a 64-bit Solaris kernel. This limitation means that you cannot put a root (/) file system on a multiterabyte file system.
- These systems do not support individual files greater than 1 Tbyte.
- The maximum number of files per Tbyte of UFS file system is 1 million. This limit is intended to reduce the time it takes to check the file system with the `fsck` command.
- The maximum quota that you can set on a multiterabyte UFS file system is 2 Tbytes of 1024-byte blocks.
- Using the `fsnap` command to create a snapshot of a multiterabyte UFS file system is not currently supported.

For more information, see “What’s New in File Systems in the Solaris 10 Release?” in the *System Administration Guide: Devices and File Systems*.

Device File System (devfs)

This feature is new in the Software Express pilot program. This feature is included in the Solaris 10 3/05 release.

The `devfs` file system manages devices in the Software Express releases. Users continue to access all devices through entries in the `/dev` directory. These entries are symbolic links to entries in the `/devices` directory. The content of the `/devices` directory is now controlled by the `devfs` file system. The entries in the `/devices` directory dynamically represent the current state of accessible devices on the system. These entries require no administration.

The `devfs` file system provides the following enhancements:

- Operations in the `/devices` directory result in attaching device entries. Unused device entries are detached.
- System boot performance is increased because only device entries that are needed to boot the system are attached. New device entries are added as the devices are accessed.

For more information, see the `devfs(7FS)` man page.

Multiterabyte Disk Support With EFI Disk Labels

This multiterabyte disk support is available only for systems that run a 64-bit kernel. This feature is new in the Software Express pilot program and in the Solaris 9 4/03 release. This feature is included in the Solaris 10 3/05 release.

The Solaris 10 OS provides support for disks that are larger than 1 terabyte (Tbyte) on systems that run a 64-bit Solaris kernel.

The Extensible Firmware Interface (EFI) label provides support for physical disks and virtual disk volumes. The UFS file system is compatible with the EFI disk label, and you can create a UFS file system that is greater than 1 Tbyte. This release also includes updated disk utilities for managing disks that are greater than 1 Tbyte.

However, the SCSI driver, `ssd`, currently supports disks only up to 2 Tbytes. If you need greater disk capacity than 2 Tbytes, use a disk and storage management product such as Solaris Volume Manager to create a larger device.

For more information on using the EFI disk label, see the *System Administration Guide: Devices and File Systems*. This guide contains important information and restrictions. This information concerns using the EFI disk label with existing software products.

The Solaris Volume Manager software can also be used to manage disks that are greater than 1 Tbyte in this Solaris release. See “[Multiterabyte Volume Support in Solaris Volume Manager](#)” on page 165.

New Configuration File for the Autofs Environment

This feature is new in the Software Express pilot program. This feature is included in the Solaris 10 3/05 release.

The new configuration file for your autofs environment, `/etc/default/autofs`, provides an additional way to configure your autofs commands and autofs daemons. Now, the same specifications that you would make on the command line can be made in this new configuration file. However, unlike the specifications you would make on the command line, this file preserves your specifications, even during upgrades to your operating system. Additionally, you no longer are required to update critical startup files to ensure that the existing behavior of your autofs environment is preserved.

You can make your specifications by using the following keywords:

- `AUTOMOUNTD_ENV` permits you to assign different values to different environments. This keyword is the equivalent of the `-D` argument for `automountd`.
- `AUTOMOUNTD_NOBROWSE` turns browsing on, or turns browsing off, for all autofs mount points. This command is the equivalent of the `-n` argument for `automountd`.
- `AUTOMOUNTD_TRACE` expands each remote procedure call (RPC) and displays the expanded RPC on standard output. This keyword is the equivalent of the `-T` argument for `automountd`.
- `AUTOMOUNTD_VERBOSE` logs status messages to the console and is the equivalent of the `-v` argument for the `automountd` daemon.

- `AUTOMOUNT_TIMEOUT` sets the duration for a file system to remain idle before the file system is unmounted. This keyword is the equivalent of the `-t` argument for the `automount` command.
- `AUTOMOUNT_VERBOSE` provides notification of `autofs` mounts, unmounts, and other nonessential events. This keyword is the equivalent of the `-v` argument for `automount`.

For more information, refer to the [automount\(1M\)](#) and the [automountd\(1M\)](#) man pages.

For further information, see the *System Administration Guide: Network Services*.

X11 Windowing Enhancements

This section describes all X11 windowing features in the Solaris 10 3/05 release that are new or have been enhanced since the Solaris 9 OS was originally distributed in May 2002. To view enhancements to the Xorg server that were made in the more recent Solaris 10 7/05 release, see “Desktop Enhancements” on page 153.

Xorg X Server

This feature is new in the Solaris Express 10/04 release.

The Xorg X server from the X.Org Foundation X11R6.8 open source release is now provided for x86 systems. This server is based on the X.Org and XFree86 open source releases, which are commonly used on Linux and other platforms. This server provides higher performance and support for a wider range of hardware, including third-party drivers written for the XFree86 and Xorg servers. The Xsun server is still provided in this release for users who need features that are not provided in the Xorg server.

For more information, see the Xorg man pages. The following command locates these man pages: `man -M /usr/X11/man Xorg`.

Xfixes Xserver Extension

This feature is new in the Solaris Express 8/04 release.

The Xfixes Xserver Extension from the open source community has been added to support comprehensive cursor support for Sec. 508 accessibility.

Xscreensaver Program

This feature is new in the Software Express pilot program and in the Solaris 9 8/03 release. This feature is included in the Solaris 10 3/05 release.

The Xscreensaver program locks your monitor screen after a period of inactivity in order to prevent others from gaining access to your system session. Xscreensaver supports a variety of animation displays that can be used while the screen is locked. This program is used in the Java Desktop System.

See the `xscreensaver(1)` man page for further information.

X Event Interception Extension (XEvIE)

This feature is new in the Software Express pilot program and in the Solaris 9 4/03 release. This feature is included in the Solaris 10 3/05 release.

XEvIE is a low-level interface that intercepts all the keyboard and mouse events so that these events can be read, consumed, or modified as needed. This X extension enables better integration of assistive technologies, including technologies in the Java Desktop System.

FreeType 2.1.x

This feature is new in the Software Express pilot program and in the Solaris 9 4/03 release. This feature is included in the Solaris 10 3/05 release.

FreeType 2.1.x is an open source library that provides a simple application programming interface (API). This API can access font content in a uniform way, independently of the file format. Additionally, some format-specific APIs can be used to access special data in the font file.

Further information on FreeType is available at <http://freetype.org>.

Xserver Virtual Screen Enhancement

This feature is new in the Software Express pilot program and in the Solaris 9 12/02 release. This feature is included in the Solaris 10 3/05 release.

The Xserver Virtual Screen enhancement provides accessibility support for the Java Desktop System. This capability enables support of magnification software on systems that have just one frame buffer.

Xrender Extension

This feature is new in the Software Express pilot program and in the Solaris 9 12/02 release. This feature is included in the Solaris 10 3/05 release.

The new Xrender feature increases performance for applications, such as the StarOffice software suite, that run on the Solaris Operating System. The Xrender feature provides a modern appearance for these applications. Xrender uses hardware processing for alpha-blending and transparency effects.

Language Support Enhancements

This section describes all language support features in the Solaris 10 3/05 release that are new or have been enhanced since the Solaris 9 OS was originally distributed in May 2002.

Unicode Version 4.0 Support

Unicode Version 3.2 support was introduced in the Software Express pilot program. In the Solaris Express 8/04 release, Solaris Unicode locales support the Unicode Version 4.0. This feature is included in the Solaris 10 3/05 release.

The Unicode Version 4.0 introduces 1226 new characters over the Unicode Version 3.2. This version also includes both normative changes and informative changes as described in “The Unicode Standard 4.0” (ISBN 0-321-18578-1). The standard can be referenced online at: <http://www.unicode.org/standard/standard.html>.

The UTF-8 character representation has been also changed to a more secure form since the UTF-8 Corrigendum that was originally published in the Unicode Version 3.1 and later updated at the Unicode Version 3.2.

This feature also implements the more secure UTF-8 character representation and byte sequences in `iconv` code conversions and OS-level multibyte functions such as the following:

- `mbtowc(3C)`
- `mbstowcs(3C)`
- `mbrtowc(3C)`
- `mblen(3C)`
- `mbsrtowcs(3C)`
- `fgetwc(3C)`
- `mblen(3C)`

For further information, see the *International Language Environments Guide*.

Code Conversions for Internationalized Domain Name Support

This feature is new in the Solaris Express 10/03 release and revised in the Solaris Express 8/04 release. For Solaris 9 users, this feature is new in the Solaris 9 4/04 release.

Internationalized Domain Name (IDN) enables the use of non-English native language names as host and domain names. To use such non-English host and domain names, application developers must convert such names into ASCII Compatible Encoding (ACE) names in their applications as specified in the RFC 3490. System administrators and end users are also required to use ACE names in existing system files and applications where the networking or system administration applications do not yet support the non-English IDNs.

This feature aids in the conversion by providing the conversion API with various supported option arguments, a dedicated IDN encoding conversion utility, and `iconv` code conversions. Refer to the following man pages for more detail:

- `libidnkit(3LIB)`
- `idn_decodename(3EXT)`
- `idn_decodename2(3EXT)`
- `idn_encodename(3EXT)`
- `idnconv(1)`
- `iconv_en_US.UTF-8(5)`

For further information, see the *International Language Environments Guide*.

New `iconv` Code Conversions

This feature is new in the Solaris Express 11/04 release.

Various new `iconv` code conversions between single-byte PC and Windows code pages and various Unicode forms have been added. Also, several major Asian code pages and UCS-2LE have been added.

Details are available at Appendix A in the *International Language Environments Guide*. See also the `iconv_en_US.UTF-8(5)` man page.

New Solaris Unicode Locales

This feature is new in the Solaris Express 7/04 release. Additional locales were added in the Solaris 9 9/04 release and in the Solaris 10 3/05 release.

The following Unicode locales have been added to Solaris:

- `nl_BE.UTF-8` (Belgium Flemish)
- `nl_NL.UTF-8` (Dutch)
- `ar_SA.UTF-8` (Saudi Arabia)
- `el_GR.UTF-8` (Greece)
- `pt_PT.UTF-8` (Portugal)
- `cs_CZ.UTF-8` (Czech Republic UTF-8)
- `cs_CZ.UTF-8@euro` (Czech Republic UTF-8 + Euro)
- `hu_HU.UTF-8` (Hungary UTF-8)

The new locales are available from system login.

Locale Administrator

This feature is new in the Solaris Express 11/03 release and in the Solaris 9 9/04 release.

Locale administrator allows the user to query and configure the locales for a Solaris OS through a command-line interface. Using the `localeadm` tool, a user can display information about locale packages that are installed on the system or that reside on a particular device or directory. The user can add and remove locales on the current system on a per-region basis. For example, the user can add all locales in the Eastern European region to the current system. See the [`localeadm\(1M\)`](#) man page.

Prior to this feature's introduction, after a system was installed the user had to add or remove individual packages to change the locales on the machine. This process is prone to error because it is easy to miss or overlook packages. This tool is a supplement to the locale selection logic in the Solaris installer. The installer is still the primary application for the correct installation of Solaris locales.

Standard Type Services Framework

This feature is new in the Solaris Express 9/03 release and in the Solaris 9 4/04 release.

The Standard Type Services Framework (STSF) is a pluggable object-based architecture that allows users to access typographically sophisticated text layout and rendering. The pluggable architecture of the framework gives users the ability to use different font rasterization engines and text layout processors to achieve the desired visual representation. The pluggable architecture also manages fonts and enables application-specific fonts to be created. STSF includes both a standalone API and an X server extension to handle rendering on the server side for improved efficiency. STSF is an open source project that is sponsored by Sun Microsystems.

For more information about the project and how to use the API, see <http://stsf.sourceforge.net>.

Auto Encoding Finder

This feature is new in the Solaris Express 9/03 release.

The Auto Encoding Finder is a useful utility for global character handling. Through a general-purpose interface, Auto Encoding Finder provides an easy way to detect the encoding of particular file or string, simplifying access to various language character encodings. For example, the utility simplifies the display of web pages that do not specify encoding information. Search engines, knowledge databases, and machine translation tools might also need to detect the encoding of the language data being accessed. The Auto Encoding Finder tool simplifies this process.

See the [`auto_ef\(1\)`](#) or [`libauto_ef\(3LIB\)`](#) man pages for more detail.

SunIM-based Korean Language Engine With Auxiliary Window Support

This feature is new in the Software Express pilot program. This feature is included in the Solaris 10 3/05 release.

Korean users of the Solaris Operating System can benefit from more comprehensive keyboard input method support for the Korean language. The new Korean Language Engine with auxiliary window support offers Korean users four auxiliary windows to control and configure the Korean input method (IM).

- User-based preferences can be set within one window.
- A virtual keyboard environment can be set within another window for point-and-click selection of Korean characters.
- Within another window, users can choose the symbols that they need from special characters that are based on code points.
- Users can organize all of the windows within a special palette of control.

This IM supports three separate keyboard layouts: 2 beol sik, 3 beol sik 390, and 3 beol sik final.

Common Transliteration-based Input Method for All Indian Languages

This feature is new in the Software Express pilot program and in the Solaris 9 12/03 release. This feature is included in the Solaris 10 3/05 release.

Users who operate within any Unicode (UTF-8) locale of the Solaris Operating System can now easily and intuitively input characters from Indian regional languages. Users who interact with CDE applications, StarOffice, or Mozilla can more easily interact with Indian scripts. After selecting the transliteration-based input method (IM), users can type phonetic equivalents of Indian language scripts in English. These equivalents are then displayed in the script that is selected, and are correctly shaped and rendered with the help of an underlying layout and shaper module. As transliteration is the most commonly used input method for Indian languages, this support can greatly enhance the usability of the eight Indian scripts that are provided in the Solaris software.

Additional Keyboard Support

This feature is new in the Software Express pilot program and in the Solaris 9 4/03 release. This feature is included in the Solaris 10 3/05 release.

Software support for four additional keyboards has been added to the Solaris software: Russian keyboard Type 6, Estonian keyboard Type 6, French-Canadian keyboard Type 6, and Polish programmer's keyboard Type 5. This software support gives users in Russia, Canada, Estonia, and Poland greater flexibility for keyboard input. The support modifies standard U.S. keyboard layouts to different language needs.

Wubi Input Method

This feature is new in the Software Express pilot program and in the Solaris 9 4/03 release. This feature is included in the Solaris 10 3/05 release.

The Wubi input method (IM) is widely used in China. The encoding rule for Wubi IM is based on the radical or stroke shape of Chinese characters. Users can rapidly type Chinese characters through a standard keyboard rather than through slower, phonetic-based input methods.

Input Method Support for Indic

This feature is new in the Software Express pilot program. This feature is included in the Solaris 10 3/05 release.

Input support for Indian regional language keyboards has been added to the Solaris Operating System. Indic language users can type Indic language characters by using their preferred keyboard layouts in the Solaris Operating System.

For further information, see the *International Language Environments Guide*.

Seven Additional Indic Scripts for Support in Unicode Locales

This feature is new in the Software Express pilot program and in the Solaris 9 4/03 release. This feature is included in the Solaris 10 3/05 release.

In addition to the current support for Hindi, the following Indic scripts are now supported in the Solaris software:

- Bengali
- Gurmukhi
- Gujarati
- Tamil
- Malayalam
- Telugu
- Kannada

Speakers of these Indian regional languages have language support in the Solaris Operating System for any of the Unicode locale environments that Solaris supports.

For further information, see the *International Language Environments Guide*.

HKSCS-2001 Support in Hong Kong Locales

This feature is new in the Software Express pilot program. This feature is included in the Solaris 10 3/05 release.

HKSCS-2001 is a new version of the Hong Kong Supplementary Character Set (HKSCS). This new version adds 116 characters to the previous HKSC-1999 character set. HKSCS-2001 is supported in these Solaris Hong Kong locales: zh_HK.BIG5HK and zh_HK.UTF-8.

Additional Software

This section describes additional software features in the Solaris 10 3/05 release that are new or have been enhanced since the Solaris 9 OS was originally distributed in May 2002.

Sun Remote Services Net Connect 3.1.1

Sun Remote Services (SRS) Net Connect 3.1.1 was made available with the Solaris Express 10/04 release on the Extra Value CD. SRS Net Connect 3.1.1 enables you to self-monitor systems, create performance and trend reports, and receive automatic notifications. These features help you maximize system availability and reliability and manage potential issues.

Note – For Solaris 9 users, this feature is new in the Solaris 9 4/04 release.

For information about installation of SRS Net Connect, see the *Sun Remote Services Net Connect Activation Guide*.

Freeware Enhancements

This section describes all freeware features in the Solaris 10 3/05 release that are new or have been enhanced since the Solaris 9 OS was originally distributed in May 2002.

GCC v.3.4.3

This feature is new in the Solaris 10 3/05 release. See “GCC Version 3.4.3” on page 235.

Webmin Tool

This feature is new in the Solaris Express 11/04 release.

Webmin is a web-based system administration tool that is written in Perl. Webmin is designed to run on all UNIX and UNIX-like systems, including the Solaris OS. The tool contains many standard modules for the administration of UNIX facilities that are included in the Solaris OS. Webmin also contains modules for administration of facilities that are provided by other open source projects and by commercial vendors. Additional third-party modules are available to administer facilities and products that are not supported by the standard modules.

Webmin is open source software. Documentation is available at <http://www.webmin.com>.

Intelligent Platform Management Interface

This feature is new in the Solaris Express 11/04 release.

Intelligent Platform Management Interface (IPMI) is an industry standard that supports some aspects of “lights-out” server management. Most x86-based servers have a built-in management

controller (bmc) that enables systems to be remotely managed. For example, a system can be powered on and off, and temperatures and voltages can be measured remotely.

IPMI is open source software. Documentation is available on the web at <http://openipmi.sourceforge.net/>.

Apache Version 2

This feature is new in the Solaris Express 11/04 release.

Apache version 2 provides several improvements to Apache version 1, most notably IPv6 support. In addition, SSL/TLS has now been incorporated directly into the server, instead of being provided as a set of add-on patches and modules that must be synchronized. Apache version 1 remains the default web server in the Solaris 10 OS.

Apache version 2 is open source software. Documentation is available at <http://www.apache.org>.

BIND 9

BIND 9 is new in the Solaris Express 8/04 release. In the Solaris 10 3/05 release, the BIND version was upgraded to BIND version 9.2.4.

BIND is an open source implementation of DNS. BIND is developed by the Internet Systems Consortium (ISC). BIND allows DNS clients and applications to query DNS servers for the IPv4 and IPv6 networks. BIND includes two main components: a stub resolver API, `resolver(3resolv)`, and the DNS name server with various DNS tools.

BIND enables DNS clients to connect to IPv6 DNS servers by using IPv6 transport. BIND provides a complete DNS client-server solution for IPv6 networks.

BIND 9.2.4 is a redesign of the DNS name server and tools by the Internet Systems Consortium (ISC). The BIND version 9.2.4 nameserver and tools are available in the Solaris 10 OS.

BIND 8.x-to-BIND 9 migration information is available in the *System Administration Guide: Naming and Directory Services (DNS, NIS, and LDAP)*. Additional information and documentation about BIND 9 is also available on the ISC web site at <http://www.isc.org>. For information about IPv6 support, see the *System Administration Guide: IP Services*.

Samba Enhancements

In the Solaris Express 8/04 release, Samba, which provides file and print services to SMB/CIFS clients, has been upgraded to version 3.0.4. This version fixes several bugs. Samba was previously updated to version 3.0 in the Solaris Express 3/04 release.

Samba provides seamless file and print services to SMB/CIFS clients. Using Samba, you can share a server's disks and printers with many network clients, including the following:

- LAN Manager
- Windows for Workgroups, Windows 95, 98, and ME
- Windows NT, 2000, and XP
- Linux
- OS/2

Two UNIX clients can access file space and printers on any SMB/CIFS server. The first is an ftp-like client. The second client creates a virtual “/smb” directory that provides access to smb-shared directories and printers.

Samba 3.0 updates the SMB file and print server to further improve printing speed. In addition, Samba 3.0 provides the following enhancements:

- Adds single sign-on with Active Directory integration
- Supports international character sets
- Integrates with the LDAP
- Replaces NT primary domain controllers (PDCs) and backup domain controllers (BDCs)

For further information, see the Samba man pages. The following command locates these man pages: `man -M /usr/sfw/man samba`. See also *Using Samba, 2nd Edition* at <http://www.oreilly.com/>.

Flex 2.5.4a

This feature is new in the Solaris Express 8/04 release.

The `flex` open source lexer is an enhanced version of the `lex(1)` utility. `flex` provides the ability to generate very fast (but large) scanners and batch these scanners. `flex` has the capability to generate scanners with different globally visible variable and function names, as opposed to the `lex(1)` convention of names such as `yy_foo` or `yy_bar`.

For more information, see the man page in the `/usr/sfw/man` directory.

SIP Proxy Server

This feature is new in the Solaris Express 8/04 release.

The Session Initiation Protocol (SIP) server provides support for the deployment of VoIP/SIP services by providing a SIP proxy, registrar, and redirect server. This server is called SER. SER is RFC 3261-compliant and is available from <http://www.ipstel.org/>.

Further information is provided in the `/usr/sfw/doc/ser` directory and in the man pages in the `/usr/sfw/man` directory.

libusb 0.1.8

This feature is new in the Solaris Express 8/04 release.

The libusb library contains a set of interfaces for managing USB devices without a kernel driver.

For further information, see the `libusb` man page in the `/usr/man` directory. This man page points to more information in the `/usr/sfw/share/doc/libusb` directory.

Ghostscript 7.05

This feature is new in the Software Express pilot program and in the Solaris 9 8/03 release. This feature is included in the Solaris 10 3/05 release.

The following new freeware package is included in the Solaris Operating System.

Ghostscript 7.05 – Reads PostScript and PDF files. Displays these files on screen or converts the files to a form that is usable by many printers. See the `gs(1)` man page under `/usr/sfw/share/man`.

Note – To view license terms, attribution, and copyright statements for Ghostscript, display the license path in `/usr/sfw/share/src/<freeware name>`.

New Freeware Packages: libxml2 2.4.16 and libxslt 1.0.19

This feature is new in the Software Express pilot program and in the Solaris 9 4/03 release. This feature is included in the Solaris 10 3/05 release.

The following freeware packages are new or revised in the Software Express releases:

- `libxml2 2.4.16` – A standard for building tag-based structured documents or data
- `libxslt 1.0.19` – An XML language to define transformation for XML

Note – To view license terms, attribution, and copyright statements for these packages, display the license path in `/usr/share/src/<freeware name>/<filename>`.

ANT 1.4.1 Freeware Package

This feature is new in the Software Express pilot program and in the Solaris 9 12/02 release. This feature is included in the Solaris 10 3/05 release.

The following new freeware package is included in the Solaris Operating System: ANT 1.4.1 – This is a Jakarta ANT Java and XML-based build package.

Note – To view license terms, attribution, and copyright statements for ANT, the license path is `/usr/sfw/share/src/<freeware name>`.

Documentation Changes

This section describes major changes to the documentation in the Solaris 10 3/05 release that are new or have been enhanced since the Solaris 9 OS was originally distributed in May 2002. These changes include the following:

- “Solaris 10 Documentation DVD” on page 263
- “New Documents” on page 263
- “Document Reorganizations” on page 265
- “Documentation That Has Been Moved” on page 267
- “Documentation That Is Not Published in This Release” on page 267

To view documentation changes that are new in the Solaris 10 7/05 release, see “[Documentation Enhancements](#)” on page 157.

Solaris 10 Documentation DVD

This feature is new in the Solaris 10 3/05 release.

The Solaris 10 3/05 release includes all documentation on a new Solaris 10 Documentation DVD. This DVD replaces the CD format that was used for documentation in the Solaris 9 release.

For information about accessing Solaris 10 documentation, including instructions for using the new documentation DVD, see the *[Solaris 10 Documentation README](#)*.

New Documents

The following documents are new in the Solaris 10 product.

Device Driver Tutorial

This tutorial provides hands-on information about how to develop device drivers for the Solaris Operating System. This book includes descriptions of writing, building, installing, loading, and testing real device drivers. These instructions help you understand how drivers control devices. This book also gives an overview of the driver development environment, tools available to you to develop drivers, and techniques for avoiding some driver development problems.

Introduction to the Solaris Developer Documentation

The Solaris Operating System provides developers with numerous interfaces, frameworks, and tools to take advantage of Solaris technologies. This book provides an overview of the Solaris OS, abstracts of key documentation for Solaris developers, and links to sources of detailed information.

Solaris Dynamic Tracing Guide

The Solaris Express 1/03 release introduced the new *Solaris Dynamic Tracing Guide*. The guide includes a complete feature reference and examples for new users. See “[DTrace Dynamic Tracing Facility](#)” on page 162 for further information about DTrace.

Solaris Security for Developers Guide

The *Solaris Security for Developers Guide* describes the public application programming interfaces (API) and service provider interfaces (SPI) for the security features in the Solaris OS. This book is intended for C-language developers who want to write the following types of programs:

- Privileged applications that can override system controls
- Applications that use authentication and related security services
- Applications that need to secure network communications
- Applications that use cryptographic services
- Libraries, shared objects, and plug-ins that provide or consume security services

The book describes the following public Solaris interfaces for security:

Process Privilege Interfaces	Process privileges permit developers to enable the delegation of security overrides in privileged applications.
PAM	Pluggable authentication modules for the initial authentication of a user to a system.
GSS-API	Generic security service application program interface, for secure communication between peer applications. GSS-API provides authentication, integrity, and confidentiality protection services as well.
SASL	Simple authentication and security layer, used largely by protocols for authentication, privacy, and data integrity. SASL is intended for higher-level network-based applications.
Cryptographic Framework	A framework based on standard PKCS #11 interfaces that accommodates consumers and providers of cryptographic services.
Smart Card	A set of interfaces for developers of IFD handlers for smart card terminals.

Working examples are provided.

Solaris System Management Agent Administration Guide

The System Management Agent (SMA) is based on the open source Net-SNMP agent. This book is for administrators who want to use the System Management Agent to manage

network devices securely, and to migrate their SNMP solution from the Solstice Enterprise Agents software to the System Management Agent. This book includes a chapter that covers security topics with examples provided.

Solaris System Management Agent Developer's Guide

The System Management Agent (SMA) is based on the open source Net-SNMP agent. This book provides information for developers who want to create MIB modules to extend the functionality of the agent.

System Administration Guide: Solaris Containers-Resource Management and Solaris Zones

The Solaris Express 2/04 release introduced the new *System Administration Guide: Solaris Containers-Resource Management and Solaris Zones*. This book is for anyone responsible for administering one or more systems that run the Solaris 10 OS. The book covers resource management topics such as projects and tasks, extended accounting, resource controls, and dynamic resource pools. This book also covers virtualization with Solaris Zones. See “[Solaris Zones Software Partitioning Technology](#)” on page 178 and all descriptions in “[System Resources Enhancements](#)” on page 178.

x86 Assembly Language Reference Manual

Introduced in the Solaris Express 6/04 release, the *x86 Assembly Language Reference Manual* documents the syntax of the Solaris assembly language for x86 systems. This book is provided to help experienced assembly language programmers understand disassembled output of Solaris compilers. This book is neither an introductory book about assembly language programming nor a reference manual for the x86 architecture.

Document Reorganizations

The following documentation has been reorganized in the Solaris 10 product.

- “[Resource Management Documentation Changes](#)” on page 265
- “[System Administration Guides Reorganized](#)” on page 265
- “[Changes to the Solaris Installation Guide](#)” on page 266

Resource Management Documentation Changes

The resource management chapters are in the new volume, *System Administration Guide: Solaris Containers-Resource Management and Solaris Zones*. The Solaris 9 book, *System Administration Guide: Resource Management and Network Services*, is now *System Administration Guide: Network Services* in the Solaris 10 OS.

System Administration Guides Reorganized

The Solaris 10 system administration guides have been reorganized.

The content in the *System Administration Guide: Basic Administration* has been reduced to include only the following topics:

- Using Solaris Management Console Tools
- Managing Users Accounts and Groups

- Managing Server and Client Support
- Shutting Down and Booting a System
- Managing Software
- Managing Solaris Patches

Complex device and file system topics were moved to the *System Administration Guide: Devices and File Systems*. This new guide includes device and file system administration topics including the following:

- Managing Removable Media
- Managing Devices
- Managing Disks
- Managing File Systems
- Backing Up and Restoring File Systems

The *System Administration Guide: Advanced Administration* contains the same topics from the previous Solaris releases.

Changes to the *Solaris Installation Guide*

The contents of the *Solaris Installation Guide* have been broken into five books that cover specific topics. For information about the content of each new book, see the following descriptions.

- *Solaris 10 Installation Guide: Basic Installations*
This book describes a basic OS installation with a graphical user interface (GUI).
- *Solaris 10 Installation Guide: Network-Based Installations*
This book describes how to perform a remote Solaris installation over a local area network or a wide area network.
- *Solaris 10 Installation Guide: Solaris Live Upgrade and Upgrade Planning*
This book provides planning information for using CD or DVD media to upgrade a system to the Solaris Operating System. This book also describes how to use Solaris Live Upgrade to create and upgrade new boot environments.
- *Solaris 10 Installation Guide: Solaris Flash Archives (Creation and Installation)*
This book provides instructions for creating Solaris Flash archives and for using Solaris Flash archives to install the Solaris OS on multiple systems.
- *Solaris 10 Installation Guide: Custom JumpStart and Advanced Installations*
This book describes how to create the files and directories necessary to perform an unattended custom JumpStart installation. How to create RAID-1 volumes during a JumpStart installation is also described.

Documentation That Has Been Moved

These documents were once published as part of the Solaris 9 release. The documents are now either published elsewhere, or their content has been moved into a different document.

- *About Solaris 9 Documentation* information has been moved to the *Solaris 10 Documentation README* and to this book.
- *man pages section 3: Realtime Library Functions* information has been moved from *man pages section 3: Threads and Realtime Library Functions* to *man pages section 3: Basic Library Functions*.
- The printed *Solaris 9 Installation Release Notes* is available instead in the online *Solaris 10 Release Notes*.
- The *GSS-API Programming Guide*, which was available in previous Solaris releases, has been subsumed by the new *Solaris Security for Developers Guide*.

Documentation That Is Not Published in This Release

These documents are not published as part of the Solaris 10 product. The documents were previously published as part of the Solaris 9 release and can be found at <http://docs.sun.com>.

- All books in the *CDE Developer Collection*
- All books in the *KCMS Collection*
- *Federated Naming Server Programming Guide*
- *Solstice Enterprise Agents 1.0 User Guide*

Features by Software Express Release Date

Many of the features in the Solaris 10 releases are also included in the Software Express program. This section lists all feature titles by original release dates.

Note – For a summary of features that were introduced in the Solaris 9, Solaris 8, or Solaris 7 releases, see *What's New in the Solaris 9 Operating Environment* at <http://docs.sun.com>. This book covers Solaris 9 features, and Solaris 8 and Solaris 7 features are fully described in the appendix.

New Features in Solaris 10 1/06 Release

The Solaris 10 1/06 release introduces the following features:

- “Upgrade Support Changes for Solaris Releases” on page 141
- “Sun Update Connection, System Edition 1.0” on page 141
- “x86: GRUB Based Booting” on page 142
- “Upgrading the Solaris OS When Non-Global Zones Are Installed” on page 143
- “Enhancements to Java Desktop System Release 3” on page 145
- “New Command for `embedded_su`” on page 148
- “SunVTS 6.1” on page 150

- “New nge Driver” on page 157
- “Solaris 10 1/06 Documentation” on page 157
- “Sun Fire Documentation Changes” on page 157

The following features were introduced in Software Express releases. These features are also included in the Solaris 10 1/06 release.

New Features in Solaris Express 8/05

The Solaris Express 8/05 release introduced the following features:

- “Set SMTP to Use Transport Layer Security” on page 147
- “SCSI HBA Driver for LSI MegaRAID 320-2x SCSI RAID Controller” on page 156
- “USB CCID IFD Handler” on page 156

New Features in Solaris Express 7/05

The Solaris Express 7/05 release introduced the “x86: New prtconf Option to Display Product Names” on page 151 feature.

New Features in Solaris Express 6/05

The Solaris Express 6/05 release introduced the following features:

- “Source-Filtered Multicasting” on page 144
- “Enhancements for bge and xge Network Interfaces” on page 144
- “Large Pages Enhancement” on page 149
- “Kernel Page Relocation” on page 149
- “Support for SCSI Disks Larger Than 2 Terabytes” on page 150
- “Advanced DDI Interrupts” on page 152
- “X Client Support for XFree86 Extensions” on page 154
- “Support for Virtual USB Keyboard and Mouse Devices” on page 155
- “Memory Placement Optimization Hierarchical Lgroup Support” on page 149

New Features in Solaris Express 3/05

The Solaris Express 3/05 release introduced the “Additional Banner Page Printing Option in Solaris Print Manager” on page 151 feature.

New Features in Solaris Express 2/05

The Solaris Express 2/05 release introduced the following features:

- “Metaslot in the Cryptographic Framework” on page 148
- “IKE Enhancements” on page 148
- “Fibre-Channel HBA Port Utility” on page 151
- “Xorg Release 6.8.2” on page 154
- “X Client Support for XFree86 Extensions” on page 154
- “Support for iSCSI Devices” on page 155

New Features in Solaris 10 3/05 Release

The Solaris 10 3/05 release introduced the following features:

- “Signing ELF Objects” on page 195
- “Virtual IP Source Address Selection” on page 222
- “GCC Version 3.4.3” on page 235
- “Solaris 10 Documentation DVD” on page 263

The Solaris 10 3/05 release included enhancements to the following features:

- “Predictive Self-Healing” on page 160
- “Java Desktop System, Release 3” on page 176
- “BIND 9” on page 260
- “New Solaris Unicode Locales” on page 255

The following features were introduced in the Software Express program. These features were also included in the Solaris 10 3/05 release.

New Features in Solaris Express 11/04 Release

The Solaris Express 11/04 release introduced the following features:

- “64-bit Support on x86 Systems” on page 163
- “Kernel Selection for x86 Systems” on page 163
- “New iconv Code Conversions” on page 255
- “Webmin Tool” on page 259
- “Intelligent Platform Management Interface” on page 259
- “Apache Version 2” on page 260

The Solaris Express 11/04 release included enhancements to the following features:

- “DTrace Dynamic Tracing Facility” on page 162
- “System V IPC and Other Resource Controls” on page 180
- “New or Updated Drivers” on page 213

New Features in Solaris Express 10/04 Release

The Solaris Express 10/04 release introduced the following features:

- “Solaris Service Manager” on page 160, a new component for “Predictive Self-Healing” on page 160
- “Sun Java Web Console” on page 163
- “Multi-Owner Disk Set Support With Solaris Volume Manager for Sun Cluster” on page 164
- “Importing Remotely Replicated Disk Sets With Solaris Volume Manager” on page 165
- “Device ID Changes” on page 165
- “Package and Patch Tool Enhancements” on page 166
- “Enhancements for `pbind` and `psrset` Commands” on page 167
- “Java Desktop System, Release 3” on page 176

- “pam_ldap Changes” on page 197
- “Solaris Secure Shell Enhancements” on page 198
- “sshd Daemon and /etc/default/login” on page 199
- “New Password Options for Nonlogin and Locked Accounts” on page 199
- “-second Option to auditconfig Command Is Removed” on page 200
- “New or Updated Drivers” on page 213
- “1394 (FireWire) and Mass-Storage Support on x86 Systems” on page 215
- “Internet Printing Protocol Listener” on page 216
- “Sun Remote Services Net Connect 3.1.1” on page 259
- “Xorg X Server” on page 252

The Solaris Express 10/04 release included the following updated features:

- “Predictive Self-Healing” on page 160
- “System V IPC and Other Resource Controls” on page 180

New Features in Solaris Express 8/04 Release

The Solaris Express 8/04 release introduced the following features:

- “DHCP Event Scripting” on page 167
- “DHCP for Logical Interfaces” on page 168
- “x86: SunVTS 6.0” on page 168
- “Kernel Modular Debugger” on page 169
- “Solaris Process Accounting and Statistics Improvements” on page 169
- “Enhancements to `ls` Command” on page 170
- “Enhancements to the Java 2 Platform, Standard Edition 5” on page 183
- “OpenSSL and OpenSSL PKCS#11 Engine” on page 199
- “perzone Audit Policy” on page 200
- “Fibre Channel Connectivity for Storage Devices” on page 216
- “Stream Control Transmission Protocol” on page 223
- “Zebra Multiprotocol Routing Suite” on page 223
- “IPsec and NAT Traversal” on page 223
- “Enhancement to the `nfsmapid` Daemon” on page 224
- “sendmail Version 8.13” on page 225
- “Per-thread Mode Enhancement” on page 236
- “Perl Version 5.8.4” on page 235
- “BIND 9” on page 260
- “Samba Enhancements” on page 260
- “Flex 2.5.4a” on page 261
- “SIP Proxy Server” on page 261
- “libusb 0.1.8” on page 262

The Solaris Express 8/04 release included revisions to the following features:

- “System V IPC and Other Resource Controls” on page 180
- “USB End-User Device Support Enhancements” on page 236

- “NFS Version 4” on page 247
- “Unicode Version 4.0 Support” on page 254
- “Code Conversions for Internationalized Domain Name Support” on page 254

New Features in Solaris Express 7/04 Release

The Solaris Express 7/04 release introduced the following features:

- “New Solaris Project and Resource Management Command Functionality” on page 181
- “New Functions for Converting Strings” on page 236
- “Java Support for `pstack` Command” on page 237
- “New Solaris Unicode Locales” on page 255

New Features in Solaris Express 6/04 Release

The Solaris Express 6/04 release introduced the following features:

- “Predictive Self-Healing” on page 160
- “New `psrinfo` Option to Identify Chip Multithreading Features” on page 170
- “Enhancements to `pfiles` Tool” on page 170
- “Solaris Installation Changes Including Installation Unification” on page 186
- “Changes to PAM for the Solaris 10 OS” on page 196
- “Expanded Printer Support” on page 216
- “Using CacheFS With NFS Version 4” on page 227
- “New Mechanisms for the Solaris Cryptographic Framework” on page 237
- “Retail and Nonretail Options for Providers in Solaris Cryptographic Framework” on page 237

Also in the Solaris Express 6/04 release, important revisions were made to the following:

- “Kerberos Enhancements” on page 200
- “Sun Java System Message Queue” on page 226
- “Sun Java System Application Server” on page 227

New Features in Solaris Express 5/04 Release

The Solaris Express 5/04 release introduced the following features:

- “Common Solaris Target Disk Driver” on page 217
- “Enhancement for `vacation` Utility” on page 228
- “USB End-User Device Support Enhancements” on page 236

Also in the Solaris Express 5/04 release, revisions were made to the following:

- “DTrace Dynamic Tracing Facility” on page 162
- “Kerberos Enhancements” on page 200
- “Linkers and Libraries Updates” on page 238

New Features in Solaris Express 4/04 Release

The Solaris Express 4/04 release introduced the following features:

- “Mozilla 1.7” on page 178
- “TCP Wrappers for rpcbind” on page 202
- “CPU Performance Counters” on page 211
- “Wheel Mouse Support” on page 217
- “MILTER, sendmail's New Mail Filter API” on page 228
- “UFS Logging Enabled by Default” on page 247

New Features in Solaris Express 3/04 Release

The Solaris Express 3/04 release introduced the following features:

- “Custom JumpStart Installation Package and Patch Enhancements” on page 187
- “Layered Driver Interfaces” on page 239
- “Changes to makecontext () Function” on page 239
- “Samba Enhancements” on page 260
- “Documentation Changes” on page 263

In addition, enhancements in this release were made to “System V IPC and Other Resource Controls” on page 180.

New Features in Solaris Express 2/04 Release

The Solaris Express 2/04 release introduced the following features:

- “Solaris IP Filter” on page 170
- “Solaris Zones Software Partitioning Technology” on page 178
- “zonename Audit Token and Audit Policy Option” on page 202
- “Single UNIX Specification, Version 3” on page 240

New Features in Solaris Express 1/04 Release

The Solaris Express 1/04 release introduced the following features:

- “Expanded Disk Set Support in the Solaris Volume Manager” on page 164
- “Dynamic Resource Pools” on page 181
- “User Commands for Solaris Cryptographic Framework” on page 202
- “IKE Configuration Parameters” on page 203
- “IPv6 Advanced Sockets API” on page 229
- “Contents of /usr/lib/mail Have Moved to /etc/mail/cf” on page 229

The following features were enhanced in the Solaris Express 1/04 release:

- “Core File Content Enhancements” on page 171
- “Linkers and Libraries Updates” on page 238

New Features in Solaris Express 12/03 Release

The Solaris Express 12/03 release introduced the following features:

- “Core File Content Enhancements” on page 171
- “System Management Agent” on page 172
- “LDAP Command Changes” on page 175
- “x86: Specifying Boot Properties With `add_install_client` Command” on page 188
- “Configuring Multiple Network Interfaces During Installation” on page 188
- “Audit Time Now Reported in ISO 8601 Format” on page 203
- “Simple Authentication and Security Layer for Developers” on page 240
- “Event Ports” on page 241

New Features in Solaris Express 11/03 Release

The Solaris Express 11/03 release introduced the following features:

- “DTrace Dynamic Tracing Facility” on page 162
- “Basic Audit and Reporting Tool” on page 204
- “Added IPv6 Functionality During Solaris Installation” on page 229
- “IPv6 Temporary Addresses” on page 230
- “Locale Administrator” on page 255

New Features in Solaris Express 10/03 Release

The Solaris Express 10/03 release introduced the following features:

- “New Architecture for Network Stacks” on page 211
- “Atomic Operations” on page 241
- “Code Conversions for Internationalized Domain Name Support” on page 254

New Features in Solaris Express 9/03 Release

The Solaris Express 9/03 release introduced the following features:

- “Top-Down Volume Creation in Solaris Volume Manager” on page 166
- “Replacement of Solaris Installation Command-Line Interface” on page 189
- “IPsec and the Solaris Cryptographic Framework” on page 204
- “Solaris Support for USB Devices” on page 219
- “EHCI and OHCI Drivers” on page 221
- “sendmail Version 8.12 Uses TCP Wrappers” on page 225
- “`routadm` Command” on page 230
- “Standard Type Services Framework” on page 256
- “Auto Encoding Finder” on page 256

Feature descriptions that were revised in the Solaris Express 9/03 release include the following:

- “USB 2.0 Features” on page 218
- “USB 2.0 Devices” on page 218

- “USB Mass Storage Devices” on page 219
- “USB Driver Enhancements” on page 220
- “Sun Java System Application Server” on page 227
- “Solaris WBEM File Changes” on page 241

New Features in Software Express Pilot Program

All features that are described in this book and are not listed in the preceding sections were introduced as part of the Software Express pilot program. These features are included in the Solaris 10 3/05 release.