



Solaris Express Developer Edition What's New

Beta



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

Part No: 820-0724-07
January 2008

Copyright 2008 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. Adobe is a registered trademark of Adobe Systems, Incorporated.

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 2008 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. Adobe est une marque enregistrée de Adobe Systems, Incorporated.

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	13
What's New in Solaris Express	17
New Features in Solaris Express Developer Edition 1/08	18
Solaris Trusted Extensions Administrator's Procedures	18
Network Data Management Protocol Service	19
StarOffice 8	19
GNOME 2.20	19
Enhancements to the Solaris ZFS File System	21
x86: Support for Suspend-to-RAM "Sleep" Feature	25
x86: Virtualization Using the Sun xVM Hypervisor	25
x86: Enhanced Speedstep CPU Power Management	26
Faulty Device Retirement Feature	26
Sun StorageTek Traffic Manager	28
Improved IPsec NAT-Traversal	28
Inetd Backlog Queue Size	28
Xvnc Server and Vncviewer Client	29
64-bit SPARC: Memory Placement Optimization Support For sun4v Platforms	29
Solaris CIFS Service	29
Solaris Trusted Extensions Supports Mounting Labeled Filesystems With the NFSv3 Protocol	30
VSCAN Service	30
SPARC: Hardware Accelerated Elliptical Curve Cryptography (ECC) Support	31
Unicode-Encoding Conversion Kernel Functions	31
Unicode UTF-8 Text Preparation Kernel Functions	31
Squid Cache	31
32-bit: PHP 5	32
Ruby 1.8.6 and Rubygems 0.9.4	32

Apache 2.2 HTTP Server	32
MySQL 5.0.45	32
Perl Database Interface and PostgreSQL Driver for Perl	32
x86: GLDv3 Version bnx II Driver	33
ADMtek Fast Ethernet Driver	33
Macronix Fast Ethernet Driver	33
x86: 4965 WiFi Driver	33
x86: dmfe(7D)	33
x86: AMD-8111	34
x86: nv_sata SATA HBA Driver	34
x86: SATA ATAPI Support in AHCI Driver	34
SATA NCQ Support in AHCI Driver	34
SPARC: rtls(7D)	34
32-bit: pgAdmin III	35
GNU Libtool 1.5.22	35
VIM 7.1	35
p7zip	35
New Features in Solaris Express Developer Edition 9/07	36
Flash Player 9	36
x86: Streamlined Installer	36
Desktop Notification	36
GNOME Devhelp	36
GNOME 2.18	37
Automatic Detection and Configuration of Local USB Printers	37
Pidgin 2.0	38
GNU Automake	38
Solaris ZFS File System Features	38
uid_t and gid_t Type Changes	42
Locale Creator	43
x86: MONITOR and MWAIT CPU Idle Loop	43
IPsec and IKE Are Managed as SMF Services	44
CPU Caps	44
iSNS Support in the Solaris iSCSI Target	45
SPARC: Shared Contexts Support	45
x86: CPUID-based Cache Hierarchy Awareness	45
Flash Update Tool	46

The audit_user(4) Database Is Not Access-restricted	46
x86: Graphics Support for i945 and i965 Cards	46
rge Driver	47
x86: bnx II Ethernet Driver	47
x86: Ralink RT2500 802.11b/g Wireless Driver	47
x86: RealTek 8180L 802.11b Wireless Driver	47
x86: 3945 WiFi Driver	48
libchewing 0.3.0	48
C-URL Wrappers Library	49
Libidn - Internationalized Domain Library	49
LibGD - The Graphics Draw Library	49
Perl-Compatible Regular Expressions	49
TIDY HTML Library	50
New Features in Solaris Express Developer Edition 5/07	50
Quagga Software Routing Suite	50
DHCPv6 Client	50
nwapd Network Auto-Configuration Daemon	51
Sendmail	51
x86: Wireless WPA Supplicant	51
x86: nge Driver Updated to Support Jumbo Framework	52
Sockets Direct Protocol	52
PPD File Management Utility	52
raidctl	53
x86: stmsboot Porting	54
MPxIO Path Steering	54
x86: GRUB Extended Support for Directly Loading and Booting the unix Kernel	54
SunVTS 7.0	55
DTrace Enabled Control Block dcmd	55
Brand-Specific Handlers for zoneadm Commands	56
x86: SATA AHCI HBA Driver	56
x86: Firmware Download Capability for SATA Drives	56
x86: SCSI LOG SENSE Support Under SATA Module	57
x86: Informational Exception Control Page	57
IP Instances: LAN and VLAN Separation for Non-Global Zones	57
Improved zonecfg Procedures for Creating Containers	58
New projmod(1M) Option	59

Solaris ZFS File System Enhancements	59
Thunderbird 2.0	62
Firefox 2.0.0.3 Web Browser	62
gDesklets	63
Lightning	63
Battery Charge Monitor	63
Subversion Source Code Management System	64
GNU-diffutils	64
PostgreSQL 8.2	64
64-bit SPARC: CPU Power Management	64
Enhanced st SCSI Reservations	65
dt login Language Selection Overhaul	65
Xorg X11R7.2 Server and Drivers	66
Japanese Font Update	66
More Japanese iconv Modules for Unicode	66
Input Method Switcher Enhancement and EMEA Keyboard Layout Emulation Support	67
x86: Concurrent FPDMA READ/WRITE QUEUED Under SATA Module	67
ACM Driver of USB Communication Device Class	67
Enhanced USB EHCI Host Controller Driver	67
USCSI LUN Reset Support	68
x86: Solaris Audio Driver for ATI IXP400	68
pcwl and pcan Drivers	68
ipw and iwi Drivers	68
USB Video Class Driver	69
MPxIO Extension for Serial Attached SCSI Devices on mpt(7D)	69
HP LTO-4 Tape Drive Support	69
IBM LTO-4 Tape Drive Support	70
x86: Support for Lucent/Agere Venus Internal PCI Modems	70
SPARC: ntwdt Driver for UltraSPARC-T1 (Niagara) Systems	70
Adaptec Ultra320 SCSI Controllers	70
New Features in Solaris Express Developer Edition 2/07	71
IPsec Tunnel Reform	71
Large Send Offload	71
GNOME System Tools	71
GNOME 2.16	72

Orca	72
StarOffice 8	73
Ekiga	73
Vino	74
Solaris Live Upgrade	74
x86: Keyboard Configuration Automated	74
Upgrading the Solaris OS When Non-Global Zones Are Installed	75
Solaris Key Management Framework	77
rsync	78
sharemgr(1M) and sharectl(1M) Utilities	78
Name Service Switch Enhancements	79
SMF Enhancements to Routing Management	79
X Server DTrace Provider	79
Direct Rendering Infrastructure	80
SATA HBA Framework and Marvell Driver	80
x86: NVIDIA Accelerated Graphics Drivers	80
Adaptec aac Hardware Support	80
x86: Direct Rendering Infrastructure Porting	81
New Features in Solaris Express 12/06	81
Session Initiation Protocol Library (libsip)	81
ZFS Command History (zpool history)	81
Changes and Improvements to Removable Media Management	82
SPARC: Process Count Scalability	86
Packet Filter Hooks	87
x86: Fault Management For Next Generation AMD Opteron Processors	87
Name service switch enhancements	87
New Features in Solaris Express 11/06	88
Resource Controls to Limit Locked Physical Memory	88
x86: lx Branded Zones: Solaris Containers for Linux Applications	88
Name Service Switch and nsd Enhancements	89
Forcibly Unmount the PCFS File System	90
New Features in Solaris Express 10/06	90
System V Resource Controls for Zones	90
Internet Printing Protocol Client-Side Support	91
Selectable Use of localhost for Solaris Print Server Database Hostname	91
Single Hosts File	92

SPARC: New sysidkbd Tool Configures Your Keyboard	92
Device Naming Enhancements	93
Linker and Libraries Updates	94
New Features in Solaris Express 9/06	94
MPSS Extension to Non-ISM/DISM Anonymous Shared Memory	94
GNOME-VFS and Nautilus ACL Support	95
ZFS Option to Sort List Output	95
Improved Device in Use Error Checking	96
Non-Global Zones Now Installed With Limited Networking Configuration	97
Adobe Flash Player Plug-in for Solaris	97
New Features in Solaris Express 8/06	97
OpenSSL 0.9.8a	97
Migration of Existing EMEA, Central and South American Locales to Common Locale Data Repository	98
hostname Command	98
Zone Unique Identifier	98
Ability to Mark Zones as “Incomplete”	99
Changes to How \$TERM Value for Console is Set	99
Solaris Zones Boot Enhancements	99
x86: Solaris Audio Driver for Ultra 20 M2 Workstation	100
ld Link Editor for Object Files	100
Support for iSCSI Target Devices	101
zfs snapshot Command	101
Solaris iSCSI Initiator	102
PAPI Print Commands	102
Enhancements to fstyp Command	103
Sun Java Web Console Changes	104
New Features in Solaris Express 7/06	105
Ability to Validate a Non-Global Zone Migration Before the Migration Is Performed	105
Mailbox Size Enhancement	105
Exclusive Link-Based Failure Detection for IPMP Singleton	106
Runtime Linker Configuration File	106
Solaris Fibre Channel Host-Based Logical Unit Number Masking	107
Solaris Trusted Extensions	107
Network Services Startup	108
Gnome 2.14	108

Evince PDF and PostScript Viewer	108
Firefox 1.5	108
Thunderbird 1.5	108
Upgrading the Solaris OS When Non-Global Zones Are Installed	109
The ZFS File System	109
New Features in Solaris Express 6/06	110
IPsec Policy Bypass Option While Running the ping Command	110
Extended File Descriptor Limits For <code>stdio</code>	110
Cloning Non-Global Zones With ZFS Zonepaths and Other ZFS Enhancements	111
MDI Enabling and Disabling of a Path	111
SIGEV_THREAD Event Notification	112
x86: Uncacheable Memory Access Support	112
Predictive Self-Healing for PCI Express on x64 Systems	113
Default Desktop Session in <code>dtlogin</code>	113
New Features in Solaris Express 5/06	114
Support for Descriptive Names for Metadevices and Hot-spare Pools	114
File System Monitoring Tool (<code>fsstat</code>)	114
<code>useradd</code> Default Shell	115
ZFS Pool Import Destroyed Pools	115
Configurable Privileges for Non-Global Zones	115
SO_TIMESTAMP Socket Option	116
NFSv4 Domain Name Configurable During Installation	116
Using DTrace in a Non-Global Zone	117
64-bit SPARC: Extended Message Signaled Interrupt Support for Fire-based Platforms ..	118
IPsec Kernel Module Error Logging	118
iSCSI Logout Support	118
iSCSI MS/T Support	119
iSNS Client Support for iSCSI	119
SNIA Multipath Management API support	120
New Features in Solaris Express 4/06	121
Common Agent Container	121
Predictive Self-Healing for x64 Systems	121
Predictive Self-Healing Support for SNMP Notification	122
Java DTrace API	122
<code>mkdtemp()</code> and <code>mks temps()</code> Library Functions	123
Zone Migration in Solaris Containers Technology	123

Sun Java Web Console	123
Support for PCI Express (PCIe)	124
PostgreSQL for the Solaris OS	125
ZFS Persistent Offline	125
New Features in Solaris Express 3/06	125
32-bit: The gnome-pilot Utility	125
IP_NEXTHOP Socket Option	126
Basic Registration 1.1	126
Sun Update Connection	126
AES Counter Mode	126
x86: SATA HBA Framework Support	127
Coherent Console	127
x86: Xorg X Server Version 6.9	127
Resource Pools Service FMRIs	128
32-bit: RealPlayer for the Solaris OS	128
Zero-CountryCode Keyboard Layout Support	129
Using Compact Flash as ATA Disk	129
Zone Move and Clone Features in the Solaris Containers Technology	129
New Features in Solaris Express 2/06	130
Solaris Flash Archives	130
New fallocate Utility	130
IIIMF and Language Engines	130
32-bit: Palm Synchronization with USB Port	131
The logadm Utility	131
New Features in Solaris Express 1/06	132
ZFS Web-Based Management Feature	132
x86: PCI Express Support on x86 Systems	132
SSL Proxy Module	133
TCP_INIT_CWND Socket Option	133
x64: Watchdog Timeout Feature	133
Object Migration and Interoperability Enhancements for pkttool	134
Deimos Cryptographic Accelerator	134
HBA Drivers	134
Driver Support for New STK Tape Drive 10000 “Titanium”	135
x86: Driver Support for AMD64 Platform	135
USB-to-Serial Driver for Prolific Adapters	135

USB-to-Serial Driver for Keyspan Adapters	135
New Features in Solaris Express 12/05	136
The ZFS File System	136
IPv6 for IP Filter	140
Support for Descriptive Names in Solaris Volume Manager	141
PKCS #11 v2.20 Support in the Solaris Cryptographic Framework	141
UDP and TCP Performance Enhancement	141
Zone Rename Feature in Solaris Containers Technology	142
Advanced DDI Interrupt Framework	142
VLAN Support with xge 10Gb Ethernet Driver	142
Core Kerberos Mechanism Resync with MIT 1.4	143
SPARC: Adobe Acrobat Reader 7.0.1	143
rge Driver	143
New UTF-8 Locales	143
New Features in Solaris Express 11/05	143
Improvements to Volume Management (vol)	143
vol is Managed by the Service Management Facility (SMF)	144
UFS Utility Enhancements	145
Internet Key Exchange (IKE) Enhancement	146
cdrecord, readCD, and cdda2wav Available	146
pilot-link Software	146
New Features in Solaris Express 10/05	146
mediaLib 2.3	147
New Features in Solaris Express 9/05	147
Dynamic Interrupt Distribution	147
New Features in Solaris Express 8/05	147
Set SMTP to Use Transport Layer Security	147
Driver for IEEE-1394 Based (IIDC) Digital Camera	148
SCSI HBA Driver for LSI MegaRAID 320-2x SCSI RAID Controller	148
USB CCID IFD Handler	148
New Features in Solaris Express 7/05	149
x86: AGPgart Driver for x86 Systems	149
x86: New prt conf Option to Display Product Names	149
New Features in Solaris Express 6/05	150
x86: GRUB-Based Booting	150
Large Pages Enhancement	151

Large Pages for Kernel Memory	151
Kernel Page Relocation	152
Enhancements for bge and xge Network Interfaces	152
Source-Filtered Multicasting	152
Automatic Renewal of Kerberos Credentials	152
Enhancement to <code>ikecert certlocal</code> Command	153
Enhanced Output With the <code>metaimport</code> Command	153
Support for SCSI Disks Larger Than 2 Terabytes	153
Advanced DDI Interrupts	153
Door Interface Revisions	155
Memory Placement Optimization Hierarchical Lgroup Support	155
Support for Virtual USB Keyboard and Mouse Devices	156
X Client Support for XFree86 Extensions	156
NVIDIA CK8-04 GE Driver Support	157
New Features in Solaris Express 4/05	157
TCP Keepalive Tunables	157
New Features in Solaris Express 3/05	158
Additional Banner Page Printing Option in Solaris Print Manager	158
New Features in Solaris Express 2/05	158
Support for iSCSI Devices	158
Fibre-Channel HBA Port Utility	159
Metaslot in the Cryptographic Framework	159
IKE Enhancements	160
Xorg Release 6.8.2	160
New and Updated Drivers	160
Key Features in Solaris Express	160

Preface

Solaris Express Developer Edition What's New summarizes all features in the most current Software Express release.

Note – This Solaris™ release supports systems that use the SPARC® and x86 families of processor architectures: UltraSPARC®, SPARC64, AMD64, Pentium, and Xeon EM64T. The supported systems appear in the *Solaris 10 Hardware Compatibility List* at <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” points out specific 64-bit information about AMD64 or EM64T systems.
- “32-bit x86” points out specific 32-bit information about x86 based systems.

For supported systems, see the *Solaris 10 Hardware Compatibility List*.

Who Should Use This Book

This book provides introductory descriptions of the new Software Express features for users, developers, and system administrators.

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 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/>)

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 <code>rm filename.</code>

TABLE P-1 Typographic Conventions (Continued)

Typeface	Meaning	Example
<i>AaBbCc123</i>	Book titles, new terms, and terms to be emphasized	Read Chapter 6 in the <i>User's Guide</i> . <i>A 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	machine_name%
C shell for superuser	machine_name#
Bourne shell and Korn shell	\$
Bourne shell and Korn shell for superuser	#

What's New in Solaris Express

This document summarizes all features in the current Software Express release. The current release is the Solaris™ Express Developer Edition 1/08 release.

The Solaris Express Developer Edition (Developer release) provides a simple initial installation of the Solaris OS for your laptop. Combined with community and Sun support and training services, this release includes the tools, technologies, and platforms enables developer to create custom Solaris, Java™, and Web 2.0 applications.

The Developer release includes the following operating system, desktop, and developer tools:

- Solaris Express operating system and desktop - Includes new features from the Sun Java™ Desktop System (Java DS). Java DS is a secure and comprehensive enterprise desktop software solution that combines open source innovation from various communities such as GNOME, and Firefox. The Desktop includes the following:
 - GNOME 2.20 - The latest enhanced GNOME desktop
 - Firefox 2.0.0.3 and Thunderbird 2.0 - Current release of Mozilla's browser and email service
 - Orca - Screen reader and magnifier for the Java DS and GNOME desktop
 - Java and Gnome bindings for the GNOME Platform libraries and the Cairo 2D drawing engine - Enable GNOME and GTK+ applications to be written in Java software
 - NetBeans™ plug-ins - Used in the NetBeans IDE to create applications
 - Ekiga - An open source desktop Voice over IP (VoIP) and video conferencing application for the GNOME desktop
 - Vino - Provides the ability to remotely administer a desktop session
- Sun™ Studio Express 12/07- C, C++, Fortran compilers, IDE, and integrated tools
- NetBeans IDE 6.0 - An open-sourced IDE for Java software developers
- Java Platform, Standard Edition 6 - The OpenJDK based release of the Java platform JDK

- StarOffice™ 8 - The OpenOffice based productivity suite, including word processor, spreadsheet, and presentation tools
- Open Source Tools - Over 150 open source applications, including Perl, Python, and GCC

To learn more about the Java DS features, see <http://opensolaris.org/os/project/jds/>.

Note –

- The Developer release installation defaults to the Solaris Express Developer Edition release. This Developer release includes a set of developer tools and uses a quick installation process. The prior default was to install the Solaris Express release. The Solaris Express release does not include the developer tools set. But, this release enables you to customize your system configuration during the installation. Now, if you want to install the Solaris Express release, you must select that release in the initial installation screen.
-

New Features in Solaris Express Developer Edition 1/08

Solaris Trusted Extensions Administrator's Procedures

This system administration enhancement is new in the Developer 1/08 release.

Starting with this release, the Solaris Trusted Extensions packages are installed when the Solaris OS is installed. The `Ext raValue` 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/labeld: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 Part I, “Initial Configuration of Trusted Extensions,” in *Solaris Trusted Extensions Administrator's Procedures*.

The Developer 1/08 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*.

Network Data Management Protocol Service

This system administration enhancement is new in the Developer 1/08 release.

The Network Data Management Protocol (NDMP) is a standard for backing up data, usually to tape, from network clients. With NDMP running as a service, any NDMP-compliant data management application on the network is a client and can back up its data to the NDMP server, a Sun StorageTek NAS appliance.

StarOffice 8

This desktop tools enhancement is new in the Developer 1/08 release.

Starting with this release, StarOffice 8 has been enhanced to include a new Chart engine.

For more information about the new Chart engine, see

<http://wiki.services.openoffice.org/wiki/Chart2/Features2.3>. For more information about StarOffice, see http://www.sun.com/software/star/staroffice/whats_new.jsp.

GNOME 2.20

This desktop tools enhancement is new in the Developer 1/08 release.

GNOME 2.20 is the latest version of the multi-platform desktop environment, GNOME Desktop. GNOME 2.20 contains the following features:

- **Email client** – The email client, Evolution, contains the following features:
 - Attachment warning
 - Email notification icon in panel's notification area
 - Backup
 - “Magic Space Bar”
- **Text Editing** – Gedit, has an all-new syntax highlighting system which now supports syntax highlighting for scripting languages such as PHP and Ruby.
- **File Management** – Desktop search is integrated into the file chooser dialog. The Nautilus file manager now displays more information in the Properties window for drives, including a pie graph that displays how much space is left. In addition, you can now see the overall disk usage in the Disk Usage Analyzer utility.

- **Control Panel** – For GNOME 2.20, the control panels are reorganized slightly to reduce the number of control panels, making it easier to find what you need. For example, this release introduces Appearance control panel applet. The Theme, Background, Fonts, and Interface applets have been merged to create this new applet, simplifying the Preferences menu. In addition, some of the Accessibility preferences have been moved to a new tab in the Preferred Applications control panel.
- **Help System** – The GNOME help browser (yelp) infrastructure is modified to improve the style and layout of the help system. In addition, the colors match your current theme better. Help pages appear more quickly, as individual pages are now loaded on demand instead of the entire manual being parsed unnecessarily.
- **Right-To-Left Language Interfaces** – Right-to-left language interfaces are present for languages such as Arabic and Hebrew which are written from right to left. Users of these languages expect most user interface elements to be similarly mirrored, compared to left-to-right user interfaces.
- **GTK+** – GNOME 2.20 uses version 2.12 of the GTK+ UI toolkit API.
- **Glib** – The Glib utility library now has a `g_get_user_special_dir()` that provides the path to special folders defined by FreeDesktop.org's `xdg-user-dirs` specification and tool. For text processing, the new `GRegex` API provides regular expression string matching without the need for an additional library.
- **Glade** – Starting with this release, there are user interface and architectural improvements. For example, tool windows such as the editor, the inspector, and the palette, are now dockable.
- **Accerciser** – Accerciser is an interactive Python accessibility explorer, and a replacement of at-poke.
- **Rarian** – Rarian is a documentation meta-data library, designed as a replacement for Scrollkeeper.
- **Gnome-devel-docs** – Gnome-devel-docs is the GNOME developer documentation suite.
- **Poppler-data** – New private data is installed under `/usr/share/poppler` containing private encoding files for use with poppler.
- **GNOME Display Manager (GDM)** – GDM now has better `utmp` and `wtmp` auditing. GDM can also now use Role Based Access Control (RBAC) to control access to the Shutdown, Reboot, and Suspend features.
- **Avahi** – Some GNOME applications, such as Ekiga and Rhythmbox, provide support for service discovery and registration using Avahi. The Avahi client API can be used by all GNOME applications. The Avahi daemon makes calls to the Bonjour API and uses the Bonjour server for service discovery and registration. On Linux and FreeBSD platforms, the Avahi daemon implements the mDNS stack.

Enhancements to the Solaris ZFS File System

This sections describes new ZFS features in the Developer 1/08 release.

- **Using Cache Devices in Your ZFS Storage Pool** – In this Solaris 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
```

After cache devices are added, they gradually fill with content from main memory. Depending on the size of your cache device, it could take over an hour for them to fill. Capacity and reads can be monitored by using the `zpool iostat` command as follows:

```
# zpool iostat -v pool 5
```

Cache devices can be added or removed from the pool after the pool is created.

For more information, see `zpool(1M)` and *Solaris ZFS Administration Guide*.

- **Enhancements to the `zfs send` Command** – This release includes the following enhancements to the `zfs send` command.

- 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
```

Send all incremental snapshots between fs@snapA to fs@snapD to fs@combo.

- 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/fsclonesnap-I
.
.
# zfs receive -F pool/clone < /snaps/fsclonesnap-I
```

- 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
```

For an extended example, see “Sending and Receiving Complex ZFS Snapshot Streams” in *Solaris ZFS Administration Guide*.

- Send an incremental replication stream.

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

For an extended example, see “Sending and Receiving Complex ZFS Snapshot Streams” in *Solaris ZFS Administration Guide*.

For more information, see “Saving and Restoring ZFS Data” in *Solaris ZFS Administration Guide*.

- **ZFS Quotas and Reservations for File System Data Only** – In addition to the existing ZFS quota and reservation features, this release includes dataset quotas and reservations that do not include descendents, such as snapshots and clones, in the space consumption accounting.
 - 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
```

For more information, see “ZFS Quotas and Reservations” in *Solaris ZFS Administration Guide*.

- **ZFS File System Properties for the Solaris CIFS Service** – This release provides support for the Solaris Common Internet File System (CIFS) service. This product provides the ability to share files between Solaris and Windows or MacOS systems.

To facilitate sharing files between these systems by using the Solaris CIFS service, the following new ZFS properties are provided:

- Case sensitivity support (`casesensitivity`)
- Non-blocking mandatory locks (`nbmand`)
- SMB share support (`sharesmb`)
- Unicode normalization support (`normalization`)
- UTF-8 character set support (`utf8only`)

In addition to the ZFS properties added for supporting the Solaris CIFS software product, the `vs can` property is available for scanning ZFS files if you have a 3rd party virus scanning engine.

For more information about using these properties, see “Managing ZFS Properties” in *Solaris ZFS Administration Guide*.

For more information about the Solaris CIFS service, see the *Solaris CIFS Administration Guide*.

- **ZFS Storage Pool Properties** – This release provides ZFS pool property information. For example:

```
# zpool get all users
NAME  PROPERTY      VALUE      SOURCE
users size         16.8G     -
users used       217M      -
users available  16.5G     -
users capacity   1%        -
users altroot    -          default
users health     ONLINE    -
users guid       11063207170669925585 -
users version    8          default
users bootfs     -          default
users delegation on         default
users autoreplace off        default
users temporary  on         local
```

- 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
```

For a description of all ZFS pool properties, see “Managing ZFS Storage Pool Properties” in *Solaris ZFS Administration Guide*.

- **ZFS and File System Mirror Mounts** – In this Solaris 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
```

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

- “Solaris ZFS File System Features” on page 38
- “Solaris ZFS File System Enhancements” on page 59
- “The ZFS File System” on page 109
- “ZFS Command History (zpool history)” on page 81
- “File System Monitoring Tool (fsstat)” on page 114
- “Improved Device in Use Error Checking” on page 96

x86: Support for Suspend-to-RAM “Sleep” Feature

This system resource enhancement is new in the Developer 1/08 release.

Starting with this release, the Solaris OS includes support for the Suspend to RAM (S3) “Sleep” feature. This feature is supported on Solaris x86 based platforms that include compliant drivers, for example the Sun Ultra™ 20 M2 Workstation. For a driver to be considered compliant, the driver must support a specific feature set.

For more information, see Chapter 12, “Power Management,” in *Writing Device Drivers*.

Note – As platforms become compliant, they will be added to the list of supported workstations.

x86: Virtualization Using the Sun xVM Hypervisor

This system resource enhancement is new in the Developer 1/08 release.

The goal of virtualization is to move from managing individual datacenter components to managing pools of resources. By consolidating multiple hosts and services on a single machine, virtualization reduces costs through the sharing of hardware, infrastructure, and administration.

The Sun xVM Hypervisor is based on the work of the Xen open source community. In a running system, the Hypervisor fits between the hardware and the operating system instance. The Hypervisor can securely execute multiple virtual machines simultaneously on a single x86-compatible computer, with each virtual machine running its own operating system.

Each virtual machine instance is called a domain. There are two kinds of domains. There is one control domain, also called domain 0, or dom0. A guest operating system is called a guest domain, also referred to as domain U or domU. You can have multiple guest domains on your system.

Within Hypervisor based solutions, there are two basic types of virtualization, full virtualization and paravirtualization. The Hypervisor supports both modes. A system can have both paravirtualized and fully virtualized domains running simultaneously.

The xVM Hypervisor virtualizes the system's hardware. This means that it transparently shares and partitions the system's resources, such as CPUs, memory, and NICs, among the guest domains.

The Hypervisor runs on x64 and x86 based systems. Supported configurations include Solaris dom0, and Solaris domU, Linux domU, FreeBSD domU, and Windows domU guests. Solaris zones and branded zones can be run within a Solaris domU.

For more information, see the following:

- *System Administration Guide: Virtualization Using the Solaris Operating System*
- <http://opensolaris.org/os/community/xen>
- Man pages:
 - xVM(5)
 - virsh(1M)
 - virt-install(1M)
 - xenconsole(1M)
 - xend(1M)
 - xenstored(1M)
 - xentop(1M)
 - xm(1M)

x86: Enhanced Speedstep CPU Power Management

This device management enhancement is new in the Developer 1/08 release.

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

For more information on how to enable Solaris CPU power management, see the `power.conf(4)` man page.

Faulty Device Retirement Feature

This device management enhancement is new in the Developer 1/08 release.

Starting with this 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, or 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 faulted 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
.
.
.
```

Sun StorageTek Traffic Manager

This device management enhancement is new in the Developer 1/08 release.

Starting with this release, the mechanisms described in `scsi_vhci(7D)` to override the autoconfiguration behavior have changed. Existing customization will be converted to the new mechanism on upgrade.

For more information, see the `scsi_vhci(7D)` man page and *Solaris SAN Configuration and Multipathing Guide*.

Improved IPsec NAT-Traversal

This networking enhancement is new in the Developer 1/08 release.

Starting with this release, IPsec Key Management applications can now enable or disable NAT-Traversal through a UDP socket option, and enable the correct PF_KEY extensions on their IPsec Security Associations.

Inetd Backlog Queue Size

This networking enhancement is new in the Developer 1/08 release.

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` using which the queue size can 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:

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

- To change the value for a specific service:

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

- To change the value globally:

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

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

Xvnc Server and Vncviewer Client

This X11 windowing enhancement is new in the Developer 1/08 release.

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 Developer 1/08 release now includes Xvnc, an X server based on the open source releases from the RealVNC project and X.Org Foundation, that displays to a RFB protocol client over the network, without requiring an existing X server session displayed on local video hardware. This release also includes the RealVNC vncviewer RFB client to connect to remote VNC servers, and several associated programs for managing these.

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

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

This system performance enhancement is new in the Developer 1/08 release.

Memory Placement Optimization (MPO) enables operating systems to allocate memory local to the core where the threads or processes are executing. 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 the platforms where multiple sockets with memory access latency differences exist. The MPO feature enhances the performance of various applications by enabling the OS to allocate memory local to the nodes.

Solaris CIFS Service

This file system enhancement is new in the Developer 1/08 release.

The Solaris CIFS service provides a native, well-integrated CIFS service to support Windows, MacOS, and other CIFS clients. This service offers ubiquitous access to files that are shared between CIFS and NFS clients. The Solaris CIFS server can act as a member server within an Active Directory domain. The Solaris CIFS service provides file system access to Windows and Mac OS clients through CIFS shares with support for both local and Active Directory domain authentication.

Similar to NFS, CIFS provides network file system services. CIFS also provides services, such as network transport for sub-protocols like named pipes, MS-RPC services, and interfaces to core Windows functionality.

For more information, see the following:

- *Solaris CIFS Administration Guide*
- `smbadm(1M)`, `smbd(1M)`, `smbstat(1M)`, and `smbautohome(4)`, `smbd(1M)`, and `pam_smb_passwd(5)` man pages.

Solaris Trusted Extensions Supports Mounting Labeled Filesystems With the NFSv3 Protocol

This security enhancement is new in the Developer 1/08 release.

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 in using TCP as an underlying transport protocol for NFS. However, users cannot choose 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.

VSCAN Service

This security enhancement is new in the Developer 1/08 release.

The Solaris OS now supports integrated virus scanning of ZFS-resident files by using the ICAP protocol to send candidate files to external third-party, off-the-shelf virus scanning products.

For more information, see the following:

- *System Administration Guide: Security Services*
- `vscanadm(1M)` and `vs cand(1M)` man pages

SPARC: Hardware Accelerated Elliptical Curve Cryptography (ECC) Support

This security enhancement is new in the Developer 1/08 release.

The UltraSPARC-T2 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 and OpenSSL users.

Unicode-Encoding Conversion Kernel Functions

This kernel functions enhancement is new in the Developer 1/08 release.

Starting with this release, a set of Unicode-encoding conversion kernel and user land functions is available for the UTF-8, UTF-16, and UTF-32 Unicode encodings. Big-endian and little-endian variations of the encodings and Byte Order Mark processing are also supported.

For more information, see the `uconv_u16tou32(9F)` and `uconv_u16tou32(3C)` man pages.

Unicode UTF-8 Text Preparation Kernel Functions

This kernel functions enhancement is new in the Developer 1/08 release.

This feature introduces a new set of kernel and user land functions that can be used to perform Unicode Normalizations and Unicode simple-case conversions on UTF-8 text. There are also functions for UTF-8 string comparison and validation with various options.

For more information, see the following man pages:

- `u8_strcmp(3C)`
- `u8_strcmp(9F)`
- `u8_textprep_str(3C)`
- `u8_textprep_str(9F)`
- `u8_validate(3C)`
- `u8_validate(9F)`

Squid Cache

This Web Stack enhancement is new in the Developer 1/08 release.

Squid is a fully-featured HTTP/1.0 proxy. Squid offers a rich access control, authorization and logging environment to develop web proxy and content serving applications.

For more information, see <http://www.squid-cache.org/Versions/v2/2.6/cfgman/>.

32-bit: PHP 5

This Web Stack enhancement is new in the Developer 1/08 release.

Starting with this release, the Solaris OS includes PHP 5. PHP Hypertext Preprocessor is a popular scripting language for web application development.

For more information, see <http://www.php.net/>.

Ruby 1.8.6 and Rubygems 0.9.4

This Web Stack enhancement is new in the Developer 1/08 release.

Starting with this release, the Ruby programming language, certain extensions, the Rails application framework, and the Rubygems package management system are supported.

For more information, see the following:

- <http://www.ruby-lang.org> for information on the Ruby programming language.
- <http://www.rubygems.org> for information on Rubygems.

Apache 2.2 HTTP Server

This Web Stack enhancement is new in the Developer 1/08 release.

Starting with this release, the Solaris OS includes the Apache 2.2.6 HTTP server. The Apache server supports multiple MPMs, PHP, prefork, and worker.

For more information, see <http://httpd.apache.org/docs/2.2/>.

MySQL 5.0.45

This Web Stack enhancement is new in the Developer 1/08 release.

Starting with this release, the Solaris OS includes the MySQL 5.0.45 Relational Database Management Systems.

Perl Database Interface and PostgreSQL Driver for Perl

This additional software enhancement is new in the Developer 1/08 release.

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

For more information, see the following:

- <http://dbi.perl.org>
- <http://search.cpan.org/~dbdpg/DBD-Pg-1.49/Pg.pm>

x86: GLDv3 Version bnx II Driver

This driver enhancement is new in the Developer 1/08 release.

The Broadcom NetXtreme (bnx) II Ethernet driver is converted to GLDv3. This conversion includes some features in GLDv3 that are useful for systems based on bnx(7d) such as full support for VLANs and 802.3 link aggregation. This is also useful for additional stack features such as IP instances.

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

ADMtek Fast Ethernet Driver

This driver enhancement is new in the Developer 1/08 release.

Starting with this release, the afe(7D) driver is introduced. The afe(7D) supports network interfaces based on ADMtek Centaur and Comet chips.

Macronix Fast Ethernet Driver

This driver enhancement is new in the Developer 1/08 release.

Starting with this release, the mxfe(7D) driver is introduced. The mxfe(7D) supports 10/100 ethernet devices based on the Macronix 98715 controller.

x86: 4965 WiFi Driver

This driver enhancement is new in the Developer 1/08 release.

The new 4965 WiFi driver supports the Intel Centrino 4965 WiFi chip. The new driver is useful to laptop users with the 4965 chip in them.

x86: dmfe(7D)

This driver enhancement is new in the Developer 1/08 release.

Starting with this release, the `dmfe(7D)` driver for Davicom 10/100 Fast Ethernet devices has been updated to support x86 platforms.

x86: AMD-8111

This driver enhancement is new in the Developer 1/08 release.

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

x86: `nv_sata` SATA HBA Driver

This driver enhancement is new in the Developer 1/08 release.

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

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

x86: SATA ATAPI Support in AHCI Driver

This driver enhancement is new in the Developer 1/08 release.

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

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

SATA NCQ Support in AHCI Driver

This driver enhancement is new in the Developer 1/08 release.

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

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

SPARC: `rtls(7D)`

This driver enhancement is new in the Developer 1/08 release.

Starting with this release, the `rtls(7D)` Ethernet is updated to support SPARC platforms. For more information, see the `rtls(7D)` man page.

32-bit: pgAdmin III

This freeware enhancement is new in the Developer 1/08 release.

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 develop complex databases.

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

GNU Libtool 1.5.22

This freeware enhancement is new in the Developer 1/08 release.

GNU Libtool is a script that enables package developers to provide generic shared library support. Libtool is used by developers who are working on software that has already adopted it. It is usually used in conjunction with the other GNU auto tools, Automake and Autoconf.

VIM 7.1

This freeware enhancement is new in the Developer 1/08 release.

Vi IMproved (VIM) is a popular clone of Visual Editor (`vi`). VIM is more full-featured than the SystemV `vi` editor in `/usr/bin/vi`.

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

p7zip

This freeware enhancement is new in the Developer 1/08 release.

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/>.

New Features in Solaris Express Developer Edition 9/07

Flash Player 9

This desktop tools enhancement is new in the Developer 9/07 release.

Starting with this release, the Solaris OS includes the Adobe Flash Player 9. For more information on this Flash Player, see

<http://www.adobe.com/products/flashplayer/productinfo/features/>.

x86: Streamlined Installer

This desktop tools enhancement is new in the Developer 9/07 release.

The Developer 9/07 release provides a streamlined installation experience for x86 laptop users. A new graphical user interface guides you through a user-friendly installation of the latest Solaris Express Developer Edition release, including the most current developer tools.

For more information about this feature and step-by-step procedures, see the Solaris Express Developer Edition Installation Guide: Laptop Installations.

Desktop Notification

This desktop tools enhancement is new in the Developer 9/07 release.

Starting with this release, desktop notification is supported on Solaris. With the desktop notification support, users are notified when specific asynchronous events occur, for example, when a removable disk is plugged into the system.

For more information, see <http://www.galago-project.org/news/index.php>.

GNOME Devhelp

This desktop tools enhancement is new in the Developer 9/07 release.

Devhelp is a useful applet for developers to browse and search GNOME API references. The applet works natively with gtk-doc and is used to browse product documentation.

For more information, see the `devhelp` man pages.

GNOME 2.18

This desktop tools enhancement is new in the Developer 9/07 release.

The GNOME 2.18 enhancement has the following features:

- Improved internationalization – Pango now supports vertical text layouts in Chinese and Japanese.
- Ring chart view – The GNOME Disk Usage Analyzer has a new ring chart view.
- Evince – This feature now supports thumbnails for PostScript™ documents.
- Glade 3 support – Glade 2 is upgraded to Glade 3 for faster performance.
- Eye of GNOME – If you rotate your photos in the camera, they stay rotated when viewed with the Eye of GNOME image viewer.
- Additional spell check support – Aspell is replaced by Enchant. Enchant shares dictionaries with Thunderbird and Firefox that enable Evolution and Gedit to support spell check for more languages.
- Tracking recently opened files – GNOME now locates recently opened files so that you do not lose track of your work.
- Deskbar applet search for new distractions – The deskbar applet provides an omnipresent, versatile search interface. Searches are handled by a series of plug-ins or handlers. Users can type keywords into the deskbar applet text box and the deskbar generates a list of search results. The search results include local files, applications, and web content similar to the Yahoo handler or the Mozilla handler. The search result depends on the handlers that are enabled by users.
- Additional plug-in support – New rhythmbox plug-ins `artdisplay` or `coverart`, `mmkeys`, and `visualizer` are now available. The `artdisplay` rhythmbox plug-in is used to gather and display album covers from the Internet. The `mmkeys` plug-in controls rhythmbox by using shortcut keys. The `visualizer` plug-in provides an enhanced user experience.
- Online games – Chess with a 3-D look and Sudoku are available.

Automatic Detection and Configuration of Local USB Printers

This desktop tools enhancement is new in the Developer 9/07 release.

The Developer 9/07 release includes desktop support for the Automatic Printing Configuration in the OpenSolaris Presto Project. The desktop support includes a GNOME applet, and configuration and preferences applet, `capplet`. The applet enables notification and processing of specific printer events such as hot-plug or error events on the GNOME panel. The `capplet`

enables controlling the applet behavior. In this release, the desktop support is in combination with the HAL USB printer support which enables automatically detecting and configuring local USB printer support.

For more information, see “Support for Automatic Printer Discovery and Configuration in the GNOME Desktop Environment” in *System Administration Guide: Solaris Printing*.

Pidgin 2.0

This desktop tools enhancement is new in the Developer 9/07 release.

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.
- Overall improvements to AIM and ICQ file transfers.
- Improvements to the 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.

GNU Automake

This developer tools enhancement is new in the Developer 9/07 release.

Starting with this release, GNU Automake 1.9.6 and GNU Automake 1.10 have been integrated into the Solaris OS. GNU Automake is a tool which is used for generating makefiles and is commonly used by developers working on open-source projects.

For more information, see <http://sourceware.org/automake/>

Solaris ZFS File System Features

This section describes new ZFS features in the Solaris Express Developer Edition 9/07 release.

- **ZFS command history enhancements (zpool history)** – The `zpool history` command has been enhanced with new options that provide ZFS file system event logging and a long format that includes the user name, the hostname, and the zone in which the operation was performed.

For example, the `zpool history -i` option provides `zpool` command events and `zfs` command events.

```
# zpool history -i users
History for 'users':
2007-04-26.12:44:02 zpool create users mirror c0t8d0 c0t9d0 c0t10d0
2007-04-26.12:46:13 zfs create users/home
2007-04-26.12:46:18 zfs create users/home/markm
2007-04-26.12:46:23 zfs create users/home/marks
2007-04-26.12:46:30 zfs create users/home/neil
2007-04-26.12:47:15 zfs snapshot -r users/home@yesterday
2007-04-26.12:54:50 zfs snapshot -r users/home@today
2007-04-26.13:29:13 zfs create users/snapshots
```

The `zpool history -l` option provides a long format. For example:

```
# zpool history -l tank
History for 'tank':
2007-07-19.10:55:13 zpool create tank mirror c0t1d0 c0t11d0
[user root on neo:global]
2007-07-19.10:55:19 zfs create tank/cindys
[user root on neo:global]
2007-07-19.10:55:49 zfs allow cindys create,destroy,mount,snapshot tank/cindys
[user root on neo:global]
2007-07-19.10:56:24 zfs create tank/cindys/data
[user cindys on neo:global]
```

For more information, see `zpool(1M)`.

- Upgrading ZFS File Systems (zfs upgrade)** – Starting with this release, the `zfs upgrade` command is included to provide future file system enhancements to existing ZFS file systems. 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
```

However, no new ZFS file system upgrade features are provided in this release.

- **ZFS delegated administration** – Starting with this release, you can delegate fine-grained 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.

The following example shows how to set permissions so that user `cindys` can create, destroy, mount and take snapshots on `tank/cindys`. The permissions on `tank/cindys` are also displayed.

```
# zfs allow cindys create,destroy,mount,snapshot tank/cindys
# zfs allow tank/cindys
-----
      Local+Descendent permissions on (tank/cindys)
      user cindys create,destroy,mount,snapshot
-----
```

Because the `tank/cindys` mount point permission is set to 755 by default, user `cindys` will be unable to mount file systems under `tank/cindys`. Set an ACL similar to the following syntax to provide mount point access.

```
# chmod A+user:cindys:add_subdirectory:allow /tank/cindys
```

You can modify the ability to use ZFS 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.

For more information, see Chapter 9, “ZFS Delegated Administration,” in *Solaris ZFS Administration Guide*.

- **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 might be possible by using separate intent log devices in your ZFS storage pool, such as with NVRAM or a dedicated disk.

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

You can set up separate ZFS logging devices in the following ways:

- When the ZFS storage pool is created or after the pool is created.

- You can attach a log device to an existing log device to create a mirrored log device. This operation is identical to attaching a device in a unmirrored storage pool.

For examples on setting up log devices, see “Creating a ZFS Storage Pool with Log Devices” in *Solaris ZFS Administration Guide* and “Adding Devices to a Storage Pool” in *Solaris ZFS Administration Guide*.

For information about whether using separate ZFS logging devices is appropriate for your environment, see “Setting Up Separate ZFS Logging Devices” in *Solaris ZFS Administration Guide*.

- **Creating intermediate ZFS datasets** – You can use the `-p` option with the `zfs create`, `zfs clone`, and `zfs rename` commands to quickly create a non-existent 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 hotplugging enhancements** – Starting with this release, ZFS more effectively responds to devices that are removed. ZFS also provides a mechanism to automatically identify devices that are inserted with the following enhancements:
 - 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 the property is set to off, device replacement must be initiated by the administrator by using the `zpool replace` command. If the property is set to on, any new device which is found in the same physical location as a device that previously belonged to the pool, is automatically formatted and replaced. The default value for the `autoreplace` property 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 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.
- Hot spares are checked periodically to make sure they are online and available.

For more information, see `zpool(1M)`.

For more information about these ZFS file system enhancements, see the *Solaris ZFS Administration Guide*.

`uid_t` and `gid_t` Type Changes

This developer tools enhancement is new in the Developer 9/07 release.

Starting with this release, the `uid_t` and `gid_t` types have been changed from `long` (32-bit binaries) and `int` (64-bit binaries) types to the `uint32_t` type. This type is used for both 32-bit and 64-bit binaries. This change promotes compatibility with other operating systems.

As part of changing the type of `uid_t` and `gid_t`, the Solaris implementation also reserves the `uid_t` and `gid_t` values that range from 231 to 232 - 2. The value 232 - 1 continues to be reserved for use as a sentinel value. Values that are assigned in this new range are temporary assignments made by the new identity mapping service. Note that, these values do not persist over a system reboot. Therefore, applications should not write UID or GID values to persistent files in the file system. Nor should applications send these values across a network connection as a means of referring to a persistent identity. For persistence, use the corresponding identity name. Solaris utilities, such as the `tar` command, and the NFSv4 network protocol, already follow these best practices.

Because the new and old `uid_t` and `gid_t` types occupy the same amount of space, existing binaries are not affected. For example, binaries that are compiled with older definitions. In these instances, there is no need to recompile the code. Binaries, object files, as well as shared libraries including libraries that are linked against Solaris shared libraries, are also not affected.

The impact of these changes are as following:

- **For Newly-compiled C binaries** – When code is recompiled, standard development practices should catch any problems that are caused by changes in the `uid_t` and `gid_t` type. In cases where the change of sign might have implications for the semantics of the code, the C compiler and, or, the lint process produce certain warnings. Developers should note that any warnings generated as a result of this change should be inspected.
- **For C++ code** – As part of the function's symbol name, C++ compilers employ a technique that is called “name mangling”. This technique encodes the primitive type names that are used in function call signatures. The change in the `uid_t` and `gid_t` type results in different name mangling for C++ functions and objects.

As with C binaries, C++ binaries and libraries continue to function as before. The exception is when libraries contain interfaces that use `uid_t` and `gid_t`.

The Solaris OS does not expose C++ interfaces of this nature. Therefore, no incompatibilities with Solaris libraries are expected.

For details on the consequences of this change, including warnings, see “Changes to `uid_t` and `gid_t` Types in the Developer 9/07 Release” in *Solaris Express Developer Edition Release Notes*.

Locale Creator

This language support enhancement is new in the Developer 9/07 release.

Locale Creator is a command line and Graphical User Interface tool which enables users to create and customize Solaris locales. This tool enables users to create installable Solaris packages containing customized locale data of a specific locale. Once the created package has been installed, the user has a fully working locale available on the system. For more information, see:

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

x86: MONITOR and MWAIT CPU Idle Loop

This kernel functions enhancement is new in the Developer 9/07 release.

The Solaris OS uses the SSE3 MONITOR and MWAIT instructions in 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”. MWAIT halts the processor until the address previously specified with MONITOR is accessed. With the new idle loop a processor only has to write to memory to wake up a halted processor.

IPsec and IKE Are Managed as SMF Services

This networking enhancement is new in the Developer 9/07 release.

Starting with this release, the service management facility (SMF) manages IPsec and IKE as a set of services:

- `svc:/network/ipsec/policy:default`
- `svc:/network/ipsec/ipsecalgs:default`
- `svc:/network/ipsec/manual-key:default`
- `svc:/network/ipsec/ike:default`

By default, the `policy` and `ipsecalgs` services are enabled, and the `manual-key` and `ike` services are disabled. For traffic to be protected by IPsec, configure either manual keying or IKE, and populate the `/etc/inet/ipsecinit.conf` file. This enables the appropriate key management service before refreshing the policy service. For more information, see Chapter 19, “IP Security Architecture (Overview),” in *System Administration Guide: IP Services*.

The Network IPsec Management profile is added for managing IPsec in a role. Also, the `ipseconf` and `ipseckey` commands can check the syntax of their respective configuration files. For more information, see the `ipseconf(1M)` and `ipseckey(1M)` man pages.

CPU Caps

This system resources enhancement is new in the Developer 9/07 release.

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: Virtualization Using the Solaris Operating System*

iSNS Support in the Solaris iSCSI Target

This device management enhancement is new in the Developer 9/07 release.

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.

Currently, the Solaris iSCSI software does not include native iSNS server support, but in this Solaris release, you can use the `iscsictadm` command to add access to an existing iSNS server to automatically discover the iSCSI devices in your network. The iSNS server can be specified by hostname or IP address. After you add the iSNS server information, you will need to enable access to the server.

For more information, see `iscsictadm(1M)` and Chapter 14, “Configuring Solaris iSCSI Targets and Initiators (Tasks),” in *System Administration Guide: Devices and File Systems*.

SPARC: Shared Contexts Support

This system performance enhancement is new in the Developer 9/07 release.

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 are because the data at a particular shared memory the address in different processes may really be identical, but the context number associated with each process is different. Therefore, the MMU hardware will not be able to recognize a match. This results in mappings being unnecessarily evicted from the MMU translation cache, Translation Lookaside Buffer (TLB), to be replaced by identical mappings with a different context number.

The UltraSPARC T2 (Niagara 2) system has an additional “shared” context, which is a hardware feature which can be used to prevent the inefficiency in handling shared memory. When the TLB is searched 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 which use Dynamic Intimate Shared Memory (DISM). In this case the process text segment and DISM segments mapped at the same virtual address and with the same permissions for each process will use the shared context.

x86: CPUID-based Cache Hierarchy Awareness

This system performance enhancement is new in the Developer 9/07 release.

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

Flash Update Tool

This system administration enhancement is new in the Developer 9/07 release.

`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.

The `audit_user(4)` Database Is Not Access-restricted

This security enhancement is new in the Developer 9/07 release.

The `audit_user(4)` database is no longer an access-restricted database. Existing Network Information Service Plus (NIS+) name servers can be updated with the command:

```
# nischmod nw+r audit_user
```

Ensure the following conditions:

- The SolarisAuditUser object class entries should be publicly readable. Existing LDAP directory name servers should not modify the SolarisAuditUser object class entries, SolarisAuditAlways and SolarisAuditNever acis, from the default, `idsconfig(1M)`.
- If the `/usr/lib/ldap/idsconfig` file has customized acis for the SolarisAuditAlways and SolarisAuditNever object class entries, the customized acis should allow anonymous read access.

Creating a new NIS name server using the installed `ypmake(1M)` creates an unrestricted `audit_user` database. Similarly, creating a new NIS+ name server using `nissetup(1M)` also creates an unrestricted `audit_user` database.

For more information, see the following:

- `audit_user(4)` man page

x86: Graphics Support for i945 and i965 Cards

This driver enhancement is new in the Developer 9/07 release.

Starting with this release, 2-D and accelerated 3-D graphics are supported for Intel integrated i945 and i965 series graphics cards. This support includes the kernel agpart and Direct Rendering Manager (DRM) drivers, user level Xorg DDX drivers and corresponding Mesa 3-D drivers.

The graphics support feature enables high resolution and accelerated 3-D rendering on the i945 and i965 graphics cards.

rge Driver

This driver enhancement is new in the Developer 9/07 release.

Starting with this release, support is provided for the RTL8110SC/RTL8169SC chipset in the Realtek Gigabit Ethernet driver rge.

x86: bnx II Ethernet Driver

This driver enhancement is new in the Developer 9/07 release.

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.

x86: Ralink RT2500 802.11b/g Wireless Driver

This driver enhancement is new in the Developer 9/07 release.

Starting with this release, a new wireless driver `ral` is integrated into the Solaris OS. The new driver provides support for the Ralink RT2500 802.11b/g chipset. The driver increases the coverage of WiFi support in Solaris and enhances user experience on Solaris WiFi.

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

x86: RealTek 8180L 802.11b Wireless Driver

This driver enhancement is new in the Developer 9/07 release.

Starting with this release, `rtw(7D)` is the driver for RealTek 8180L 802.11b PCMCIA wireless adapter. The driver increases the coverage of WiFi support in Solaris and enhances user-experience on Solaris WiFi.

For more information, see the following:

- `rtw(7D)` man page
- <http://opensolaris.org/os/community/laptop/>

x86: 3945 WiFi Driver

This driver enhancement is new in the Developer 9/07 release.

The new 3945 WiFi driver supports the Intel Centrino 3945 WiFi chip. The new driver is useful to laptop users with the 3945 chip in them.

libchewing 0.3.0

This language support enhancement is new in the Developer 9/07 release.

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:

- Incompatibility with API/ABI.
- UTF-8 based language engine core for common Unicode environment.
- UTF-8 based language engine core for common Unicode environment.
- UTF-8 based language engine core for common Unicode environment.
- Zuin fixes and symbol improvements.
- Some fixed internal memory management bugs and leaks.
- New binary-form of user hash data to speed up loading and solving hash data corruption.
- Better calculating 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 which recomputes 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*.

C-URL Wrappers Library

This freeware enhancement is new in the Developer 9/07 release.

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

This freeware enhancement is new in the Developer 9/07 release.

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:

- <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

This freeware enhancement is new in the Developer 9/07 release.

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/>.

Perl-Compatible Regular Expressions

This freeware enhancement is new in the Developer 9/07 release.

PCREs enable programmatic access to Perl-Compatible Regular Expressions. Prior to PCRE, Perl regular expressions were only available through Perl.

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

TIDY HTML Library

This freeware enhancement is new in the Developer 9/07 release.

TIDY is a HTML parser. It 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/>.

New Features in Solaris Express Developer Edition 5/07

Quagga Software Routing Suite

This networking enhancement is new in the Developer 5/07 release.

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 'routeadm'.

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

This networking enhancement is new in the Developer 5/07 release.

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:

- `dhcagent(1M)`
- `in.ndpd(1M)`
- `ifconfig(1M)`
- `ndpd.conf(4)`
- `dhcpinf(1)`

nwamd Network Auto-Configuration Daemon

This networking enhancement is new in the Developer 5/07 release.

The Developer 5/07 release booting process runs the `nwamd` daemon. This daemon implements an alternate instance of the SMF service, `svc:/network/physical`, which enables automated network configuration with minimal intervention.

The `nwamd` daemon monitors the Ethernet port and automatically enables DHCP on the appropriate IP interface. If no cable is plugged into a wired network, the `nwamd` daemon conducts a wireless scan and sends queries to the user for a WiFi access point to connect to.

Laptop users do not need to spend extensive amounts of time manually configuring the interfaces on their systems. Automatic configuration also aids system administrators, who can reconfigure network addresses with minimal intervention.

The OpenSolaris Network Auto-Magic Phase 0 page and `nwamd` man page contain further details, including instructions for turning off the `nwamd` daemon, if preferred. For more information and a link to the `nwamd(1M)` man page, see <http://www.opensolaris.org/os/project/nwam/phase0/>.

Sendmail

This networking enhancement is new in the Developer 5/07 release.

There are some minor new options to `sendmail(1M)` and some new configuration macros described in `/etc/mail/cf/README`. The chief of these macros are the following three new `FEATURE()` macros, all of which are helpful in blocking unwanted spam:

- `badmx`
- `require_rdns`
- `block_bad_helo`

For information about `sendmail`, see the `sendmail(1M)` man page.

x86: Wireless WPA Supplicant

This networking enhancement is new in the Developer 5/07 release.

The Solaris WiFi driver has been enhanced to support WiFi Protected Access (WPA) personal mode. Starting with this release, a new service, `network/wpa`, is introduced. The `nwam(1M)`, `dladm(1M)`, `net80211(5)`, and `ath(7D)` drivers are enhanced to support the WPA personal mode. Users can access the WPA wireless network by using `dladm` or `nwam`:

```
$ dladm create-secobj -c wpa mykey /* create a wpa psk specified by your AP */
```

```
$ dladm connect-wifi -e essid -k mykey
```

x86: nge Driver Updated to Support Jumbo Framework

This networking enhancement is new in the Developer 5/07 release.

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.

Sockets Direct Protocol

This networking enhancement is new in the Developer 5/07 release.

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 which is very similar to TCP.

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

PPD File Management Utility

This system administration enhancement is new in the Developer 5/07 release.

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. 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. The `print/ppd-cache-update` service is enabled by default.

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*

raidctl

This system administration enhancement is new in the Developer 5/07 release.

`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:

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

x86: stmsboot Porting

This system administration enhancement is new in the Developer 5/07 release.

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>.

MPxIO Path Steering

This system administration enhancement is new in the Developer 5/07 release.

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 allows network administrators to run diagnostic commands through a specified path.

x86: GRUB Extended Support for Directly Loading and Booting the `unix` Kernel

This system administration enhancement is new in the Developer 5/07 release.

Starting with this release, changes have been made to GRUB that enable the boot loader to directly load and boot the `unix` kernel. The GRUB `multiboot` module is no longer used. This implementation integrates the previous multiboot functionality directly into the platform-specific kernel module. These changes reduce the time, as well as memory requirements, that are needed to boot the Solaris OS.

Other additions and modifications to GRUB based booting include:

- The addition of two new keywords, `kernel$` and `module$`.

These keywords assist in the creation of menu .lst file entries that work with both 32-bit and 64-bit systems.

- The bootadm command has also been modified to create GRUB menu .lst file entries that contain the platform-specific unix module, which is now loaded directly by GRUB.

During a software upgrade, the bootadm command converts any multiboot entries in the menu .lst file to directly refer to the unix kernel module. For more information, see the boot(1M) and bootadm(1M) man pages.

For step-by-step instructions, see Chapter 11, “Modifying Solaris Boot Behavior (Tasks),” in *System Administration Guide: Basic Administration* and “Booting an x86 Based System by Using GRUB (Task Map)” in *System Administration Guide: Basic Administration*.

SunVTS 7.0

This system administration enhancement is new in the Developer 5/07 release.

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.

DTrace Enabled Control Block dcmd

This system administration enhancement is new in the Developer 5/07 release.

Starting with this release, the `::dt race_ecb dcmd` to the modular debugger, `mdb`, enables a user to iterate over the Enabled Control Blocks (ECBs) for a given DTrace user. This feature enables easy access to the data structures related to a user's currently enabled probes.

Brand-Specific Handlers for zoneadm Commands

This system administration enhancement is new in the Developer 5/07 release.

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: SATA AHCI HBA Driver

This system administration enhancement is new in the Developer 5/07 release.

The AHCI driver is a SATA framework-compliant HBA driver that supports various SATA HBA controllers that are compatible with the AHCI specification defined by Intel. Currently, the AHCI driver supports INTEL ICH6 and VIA vt8251 controllers, and hot-plugging functions.

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

x86: Firmware Download Capability for SATA Drives

This system administration enhancement is new in the Developer 5/07 release.

This feature provides firmware download capability for SATA drives under the SATA framework. System administrators can use the USCSI interface to update firmware on SATA disks.

x86: SCSI LOG SENSE Support Under SATA Module

This system administration enhancement is new in the Developer 5/07 release.

The SCSI LOG SENSE support feature provides the capability of obtaining the following information:

SCSI LOG SENSE PAGES 0	The available pages
0x10	Self test results
0x2f	Information Exception log
0x30	SMART READ DATA

x86: Informational Exception Control Page

This system administration enhancement is new in the Developer 5/07 release.

The SATA framework now supports the Informational Exception Control page. The Informational Exception Control page enables and disables Informational Exception reporting, which is necessary for device-health monitoring.

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

This system resources enhancement is new in the Developer 5/07 release.

IP networking can now be configured in two different ways, depending on whether the zone 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 “[x86: lx Branded Zones: Solaris Containers for Linux Applications](#)” on page 88.

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: Virtualization Using the Solaris Operating System*
For configuration information, see Chapter 17, “Non-Global Zone Configuration (Overview),” in *System Administration Guide: Virtualization Using the Solaris Operating System* and Chapter 18, “Planning and Configuring Non-Global Zones (Tasks),” in *System Administration Guide: Virtualization Using the Solaris Operating System*.
For information on feature components, see Chapter 26, “Solaris Zones Administration (Overview),” in *System Administration Guide: Virtualization Using the Solaris Operating System* and Chapter 27, “Administering Solaris Zones (Tasks),” in *System Administration Guide: Virtualization Using the Solaris Operating System*.

Improved `zonecfg` Procedures for Creating Containers

This system resources enhancement is new in the Developer 5/07 release.

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 global property names, the preferred method.
- The new `zone.max-swap` control provides swap capping for zones through the `capped-memory` resource.
- Additional methods for setting the default scheduler in a zone, including a new `scheduling-class` property, have been added.
- 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 native zones. See [“x86: `lx` Branded Zones: Solaris Containers for Linux Applications” on page 88.](#)

- The resident set size (RSS) accounting was 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: Virtualization Using the Solaris Operating System*

New `projmod(1M)` Option

This system resources enhancement is new in the Developer 5/07 release.

Use the `projmod` command with the `-A` option to apply the resource control values found 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.

Solaris ZFS File System Enhancements

This section describes new ZFS features in the Solaris Express Developer Edition 5/07 release.

- **Recursively renaming ZFS snapshots** – You can recursively rename all descendant ZFS snapshots by using the `zfs rename -r` command.
For example, snapshot a set of ZFS file systems called `users/home@today`. Then, rename all the snapshots in one command the next day, as `users/home/@yesterday`.
Snapshots are the only dataset that can be renamed recursively.
- **GZIP compression is available** – 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** – As a reliability feature, ZFS file system metadata is automatically stored multiple times across different disks, if possible. This feature is known as *ditto blocks*.

Starting with this 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.

The benefits of storing multiple copies of ZFS user data are as follows:

- Improves data retention by allowing recovery from unrecoverable block read faults, such as media faults (bit rot) for all ZFS configurations.
- Provides data protection even in the case where only a single disk is available.
- Allows you to select data protection policies on a per-file system basis, beyond the capabilities of the storage pool.
- **Improved storage pool status information** (`zpool status`) – 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 filenames from the list of displayed inodes.
- **Sharing ZFS file system enhancements** – The process of sharing file systems has been improved. Although modifying system configuration files, such as `/etc/dfs/dfstab`, is unnecessary for sharing ZFS file systems, you can use the `sharemgr` command to manage ZFS share properties. The `sharemgr` command enables you to set and manage share properties on share groups. ZFS shares are automatically designated in the `zfs share group`.

As in previous releases, you can set the ZFS `sharenfs` property on a ZFS file system to share a ZFS file system. For example:

```
# zfs set sharenfs=on tank/home
```

Or, you can use the new `sharemgr add-share` subcommand to share a ZFS file system in the `zfs share group`. For example:

```
# sharemgr add-share -s tank/data zfs
# sharemgr show -vp zfs
zfs nfs=(
  zfs/tank/data
    /tank/data
    /tank/data/1
```

```

/tank/data/2
/tank/data/3

```

Then, you can use the `sharemgr` command to manage ZFS shares. The following example shows how to use `sharemgr` to set the `nosuid` property on the shared ZFS file systems. You must preface ZFS share paths with `/zfs` designation.

```

# sharemgr set -P nfs -p nosuid=true zfs/tank/data
# sharemgr show -vp zfs
zfs nfs=(
  zfs/tank/data nfs=(nosuid="true")
  /tank/data
  /tank/data/1
  /tank/data/2
  /tank/data/3
)

```

- **ZFS and Solaris iSCSI improvements** – 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, 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** – 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** – You use the `canmount` property to specify whether a dataset can be mounted by using the `zfs mount` command.
 - **ZFS user properties** – 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** – 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

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

- **Display all ZFS file system information** – 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** – 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. Using this option might be necessary when the file system is modified between the time a rollback occurs and the receive is initiated.
- **Recursive ZFS snapshots** – 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.

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

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 2.0.0.3 Web Browser

This desktop tools enhancement is new in the Developer 5/07 release.

Firefox 2.0.0.3 focuses on user interface innovations that help users in their common browsing tasks while interacting with search, bookmarks and history. Firefox 2.0.0.3 has improvements to tabbed browsing, RSS handling, managing extensions, security and performance.

gDesklets

This desktop tools enhancement is new in the Developer 5/07 release.

gDesklets provide an advanced architecture for desktop applets. The applets placed on the desktop are meant to help the user quickly retrieve information without hindering normal activity.

For more information, see:

- <http://www.gdesklets.de>
- <http://develbook.gdesklets.de/> for a tutorial about how to create gDesklets

Lightning

This desktop tools enhancement is new in the Developer 5/07 release.

The Lightning 0.3 feature includes the following new features:

- Local calendar
- CalDAV support
- WCAP 3.0 support
- Event or task recurrence
- Meeting arrangement
- Event View- Day View, Week View, Month View
- Event or task list view
- Basic handling of events received in email
- Event or task alarms
- Calendar import and export
- Localization support

Battery Charge Monitor

This desktop tools enhancement is new in the Developer 5/07 release.

The battery charge monitor is an applet for the GNOME panel. If the battery is available on a machine, the battery charge monitor displays the status of the battery, including the remaining charge and time. If the battery is not available, the monitor shows an AC adapter icon in the system tray.

The battery charge monitor also notifies users when the battery is fully charged or is about to lose charge. The applet for the battery charge monitor ensures that customers are notified before their laptops lose power. Thus, serious loss of data can be prevented.

Subversion Source Code Management System

This developer tools enhancement is new in the Developer 5/07 release.

Starting with this release, the Subversion source code management system is included in Solaris. Subversion is a popular source code management system which is used by many open-source projects. This system is also being used in OpenSolaris development.

The Subversion feature provides Solaris users the needed tool to participate in the development of open-source projects. For more information about Subversion, see <http://subversion.tigris.org/>.

GNU-diffutils

This developer tools enhancement is new in the Developer 5/07 release.

The Developer 5/07 release includes GNU utilities for comparing and merging files. For more information, see <http://gnu.org/software/diffutils>.

PostgreSQL 8.2

This database software enhancement is new in the Developer 5/07 release.

This feature is the latest version of the PostgreSQL Open Source relational database system. PostgreSQL 8.2 for Solaris now includes Kerberos 5 support and embedded DTrace probes.

For a comprehensive list of features and enhancements in this release, see <http://www.postgresql.org/docs/current/static/release-8-2.html>. For more information about PostgreSQL, see <http://www.postgresql.org>.

64-bit SPARC: CPU Power Management

This device management enhancement is new in the Developer 5/07 release.

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.

Enhanced `st` SCSI Reservations

This device management enhancement is new in the Developer 5/07 release.

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.

`dtlogin` Language Selection Overhaul

This X11 windowing enhancement is new in the Developer 5/07 release.

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.

Xorg X11R7.2 Server and Drivers

This X11 windowing enhancement is new in the Developer 5/07 release.

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.

Japanese Font Update

This language support enhancement is new in the Developer 5/07 release.

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

More Japanese `iconv` Modules for Unicode

This language support enhancement is new in the Developer 5/07 release.

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, in addition to 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

This language support enhancement is new in the Developer 5/07 release.

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 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: Concurrent FPDMA READ/WRITE QUEUED Under SATA Module

This device driver enhancement is new in the Developer 5/07 release.

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.

ACM Driver of USB Communication Device Class

This driver enhancement is new in the Developer 5/07 release.

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

Enhanced USB EHCI Host Controller Driver

This driver enhancement is new in the Developer 5/07 release.

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 driver enhancement is new in the Developer 5/07 release.

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.

x86: Solaris Audio Driver for ATI IXP400

This driver enhancement is new in the Developer 5/07 release.

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 new Ferrari4000 model. The `audioixp` driver follows the Solaris Audio Driver Architecture (SADA) framework.

pcwl and pcan Drivers

These driver enhancements are new in the Developer 5/07 release.

The `pcwl` driver is used for Wavelan and Prism II 802.11b devices. The `pcan` driver is for Aironet 802.11b devices. These two drivers have been used by internal users through `frkit`.

ipw and iwi Drivers

These driver enhancements are new in the Developer 5/07 release.

Starting with this release, two wireless drivers are provided: the `ipw` or the Intel Pro Wireless 2100b driver and the `iwi` or the Intel Pro Wireless 2200bg/2915abg driver. These drivers increase the coverage of WiFi support in Solaris and improve the user's WiFi experience. For more information on WiFi drivers, see the Laptop community at <http://opensolaris.org/os/community/laptop/>.

USB Video Class Driver

This driver enhancement is new in the Developer 5/07 release.

The USB video class driver, `usbvc`, supports web cameras that are compliant with the USB video class specifications at <http://www.usb.org/home>. The `usbvc` driver supports the following web cameras: :

- Logitech Quickcam Ultra Vision
- Logitech Quickcam Pro 5000
- Logitech Quickcam Fusion
- Logitech Quickcam Orbit MP
- Logitech Quickcam Pro for Notebooks

The `usbvc` driver, enables a few video applications, for example, the video conference application, Ekiga.

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

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

This driver enhancement is new in the Developer 5/07 release.

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.

HP LTO-4 Tape Drive Support

This driver enhancement is new in the Developer 5/07 release.

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

IBM LTO-4 Tape Drive Support

This driver enhancement is new in the Developer 5/07 release.

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

x86: Support for Lucent/Agere Venus Internal PCI Modems

This driver enhancement is new in the Developer 5/07 release.

The Solaris serial port driver is enhanced to support internal PCI modems based on the Lucent/Agere Venus chipset. These 56-Kbyte modems appear as ordinary serial ports in Solaris.

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

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

This driver enhancement is new in the Developer 5/07 release.

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.

Adaptec Ultra320 SCSI Controllers

This driver enhancement is new in the Developer 5/07 release.

Starting with this release, the Ultra320 SCSI HBA with PCI Express interface from Adaptec can run in native MSI interrupt mode and not in legacy fix mode.

New Features in Solaris Express Developer Edition 2/07

IPsec Tunnel Reform

This networking enhancement is new in the Developer 2/07 release.

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 `ipseccnf(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*.

Large Send Offload

This networking enhancement is new in the Developer 2/07 release.

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.

GNOME System Tools

This desktop tools enhancement is new in the Developer 2/07 release.

Starting with this release, the Java DS includes the GNOME System Tools. This feature provides several system administration tools that enable the desktop users to perform basic administration of the following system's features:

- Date & Time
- Users & Groups
- Services
- Network
- Shared Folders

The GNOME System Tools can be accessed under **Launch->Administration menu**.

GNOME 2.16

This desktop tools enhancement is new in the Developer 2/07 release.

Starting with this, the Java DS includes the new Gnome 2.16 desktop. The GNOME 2.16 desktop introduces new features and enhances the usability and performance of many commonly used applications. The desktop includes these new features:

- The GTK+ 2.10 version provides several new widgets, improved themes, an enhanced file chooser, and a new printing API.
- The user's experience of P`l`ug and P`l`ay portable media and devices is greatly improved with the introduction of Freedesktop.org's HAL and related elements in the GNOME desktop environment. You can now simply insert a blank CD/DVD, or plug in a USB mass storage device, and configure the desktop interactively.
- Orca, a new screen reader and magnifier is integrated to provide enhanced accessibility support.
- Alacarte, a menu editor enables users to edit the layout of the Launch menu.
- GTK+ Authorization (Gksu) enables the running of applications as another user, or using RBAC profiles, prompting for authorization as needed.
- GNOME System Monitor panel applet and application enables the user to monitor the state of the system.
- GNOME System Log Viewer enables the viewing and monitoring of log files.
- Vino enables the remote administration of a user's desktop, allowing the administrators to see exactly what is on the user's display.
- The Java-Gnome bindings for the GNOME Platform are added in this release, enabling GNOME and GTK+ applications to be written in Java. The Java and GNOME bindings also include support for the GNOME WYSIWYG UI designer, Glade.

Orca

This desktop tools enhancement is new in the Developer 2/07 release.

Orca is a free, open-source, flexible, and extensible screen reader that provides access to the graphical desktop. Access is through user-customizable combinations of speech, braille, or magnification.

Orca works with applications and toolkits that support the Assistive Technology Service Provider Interface (AT-SPI), which is the primary assistive technology infrastructure for Solaris and Linux operating systems. Applications and toolkits that support the AT-SPI include the GNOME GTK+ toolkit, the Java platform's Swing toolkit, OpenOffice, and Mozilla.

StarOffice 8

This desktop tools enhancement is new in the Developer 2/07 release.

StarOffice 8 includes the following new features and enhancements:

- Microsoft Office compatibility
- OpenDocument as the new default file format
- Support for XForms
- Digital signatures for documents
- Resident database wizard
- Enhanced mail merge wizard
- Applications are easier to use
- Native desktop theme
- Migration tools
- Export of documents to Adobe PDF

For more information, see

http://www.sun.com/software/star/staroffice/whats_new.jsp.

Ekiga

This desktop tools enhancement is new in the Developer 2/07 release.

Ekiga is a videoconferencing and VOIP/IP-Telephony application that enables you to make audio and video calls to remote users with H.323 or SIP hardware or software (such as Microsoft Netmeeting). It supports PC-to-PC and PC-to-Phone calls.

Ekiga also supports the following features:

- Call Forwarding on busy, no answer, or always (SIP and H.323)
- Call Transfer (SIP and H.323)
- Call Hold (SIP and H.323)
- DTMFs support (SIP and H.323)
- Basic Instant Messaging (SIP)
- Text Chat (SIP and H.323)
- Possibility to register to several registrars (SIP) and gatekeepers (H.323)
- Possibility to use an outbound proxy (SIP) or a gateway (H.323)
- Message Waiting Indications (SIP)
- Audio and Video (SIP and H.323)
- STUN support (SIP and H.323)
- DTMF support
- LDAP support (Addressbook)

Vino

This desktop tools enhancement is new in the Developer 2/07 release.

Remote Desktop (vino) provides a VNC server that integrates with GNOME, supporting multiple clients to access a GNOME desktop session remotely. It enables you to export your running desktop to another computer for remote use or diagnosis.

For more information, see the `vino-preferences` and `vino-server` man pages.

Solaris Live Upgrade

This installation enhancement is new in the Developer 2/07 release.

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 Express Installation Guide: Solaris Live Upgrade and Upgrade Planning*.

x86: Keyboard Configuration Automated

Starting with this release, the `sysidkbd` tool configures your USB language and its corresponding keyboard layout on x86 systems.

With the new `sysidkbd` 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 `sysidkbd` 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 `sysidcfg` 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 `sysidcfg` file.

For more information, see the *Solaris Express Installation Guide: Network-Based Installations*.

This feature was introduced for SPARC systems in the Solaris Express 10/06 release. See [“SPARC: New `sysidkbd` Tool Configures Your Keyboard”](#) on page 92.

Upgrading the Solaris OS When Non-Global Zones Are Installed

This installation enhancement is new in the Developer 2/07 release.

Starting with this release, you can upgrade the Solaris OS when non-global zones are installed.

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. 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. 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.

Solaris Live Upgrade creates a copy of the OS on the inactive boot environment. The inactive boot environment can be upgraded or patched when non-global zones are installed. The inactive boot environment can then be booted to become the new boot environment. 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. This exception occurs under the following circumstances:
 - If in the current boot environment the `zonecfg add fs` command was used to create a separate file system for a non-global zone
 - If this separate file system resides on a shared file system. For example, `/zone/root/export`

To prevent this separate file system from being shared in the new boot environment, the `lucreate` command now enable you to specify a destination slice for a separate file system for a non-global zone. The argument to the `-m` option has a new optional field, *zonename*. This new field places the non-global zone's separate file system on a separate slice in the new boot environment.

Note – By default, any file system other than the critical file systems, that is, `root (/)`, `/usr`, and `/opt` file systems, is shared between the current and new boot environments. The `/export` file system is a shared file system. If you use the `-m` option, the non-global zone's file system is placed on a separate slice and data is not shared. The `-m` option prevents zone file systems that were created with the `zonecfg add fs` command from being shared between the boot environments. See `zonecfg(1M)` for details.

In this example, a new boot environment named `newbe` is created. The `root (/)` file system is copied to `c0t1d0s4`. All non-global zones in the current boot environment are copied to the new boot environment. The non-global zone named `zone1` contains a file system that resides on a separate slice in a shared file system such as

/zone1/root/export. To prevent this file system from being shared, the file system is copied to a separate slice on newbe, c0t1d0s1.

```
# lucreate -n newbe -m /:/dev/dsk/c0t1d0s4:ufs \
-m /export:/dev/dsk/c0t1d0s1:ufs:zone1
```

- 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.
- Comparing boot environments is enhanced. The `lucompare` command now generates a comparison of boot environments that includes the contents of any non-global zone.
- Listing file systems with the `lufsls` command is enhanced to display a list of file systems for both the global zone and the non-global zones.

For step-by-step procedures for upgrading a system with non-global zones installed or for information about the Solaris Zones partitioning technology, see the following references:

Description	For More Information
Upgrading with Solaris Live Upgrade on a system with non-global zones	Chapter 9, “Upgrading the Solaris OS on a System With Non-Global Zones Installed,” in <i>Solaris Express Installation Guide: Solaris Live Upgrade and Upgrade Planning</i>
Creating and using non-global zones	<i>System Administration Guide: Virtualization Using the Solaris Operating System</i>
Upgrading with JumpStart	<i>Solaris Express Installation Guide: Custom JumpStart and Advanced Installations</i>
Upgrading with the Solaris installation interactive GUI	<i>Solaris Express Installation Guide: Basic Installations</i>

Solaris Key Management Framework

This security enhancement is new in the Developer 2/07 release.

The `pktool` command enables the administrator to manage PKI objects in all three 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*

rsync

This system administration enhancement is new in the Developer 2/07 release.

`rsync` is an open source utility that provides fast and incremental file transfer. This utility is used by system administrators to move, copy, synchronize data both local, and over the network. The `rsync` utility uses SSH as a secure transport between machines. `rsync` can also be used as a safe remote data backup tool for both full and incremental directory comparisons and transfers.

For more information, see <http://rsync.samba.org/>.

sharemgr(1M) and sharectl(1M) Utilities

This system administration tools enhancement is new in the Developer 2/07 release.

Starting with this release, you can use two new utilities to manage file systems and file-sharing protocols:

- The `sharemgr` utility simplifies tasks related to sharing file systems. For example, when using `sharemgr` to share file systems, to set the property values for shared file systems, or perform related tasks, you do not need to use the `share`, `shareall`, or `unshare` utility. Additionally, you do not need to edit the `/etc/dfs/dfstab` file.
- The `sharectl` utility enables you to configure and manage file-sharing protocols, for example, NFS. This utility enables you to set the client and server operational properties, display property values for a specific protocol, and obtain the status of a protocol.

For more information, see the following:

- `sharemgr(1M)` man page
- `sharectl(1M)` man page
- *System Administration Guide: Network Services*

Name Service Switch Enhancements

This system administration tools enhancement is new in the Developer 2/07 release.

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.

SMF Enhancements to Routing Management

This system administration tools enhancement is new in the Developer 2/07 release.

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.

X Server DTrace Provider

Starting with the Developer 2/07 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.freeesktop.org/>.

Direct Rendering Infrastructure

Direct Rendering Infrastructure (DRI) is an open sourced software framework for coordinating the following:

- OS kernel
- X Window System
- 3D graphics hardware
- OpenGL-based client applications

DRI enables direct access to graphics hardware in the X Window System in a safe and efficient manner. DRI also enables OpenGL hardware-accelerated, 3D rendering on UNIX-like operating systems. DRI framework and the accelerated driver for Intel have been ported to Solaris. Currently, DRI has the graphics chipsets produced by Intel, ATI, Via, and 3dfx.

SATA HBA Framework and Marvell Driver

This device management enhancement is new in the Developer 2/07 release.

Starting with this release the READ/WRITE FPDMA QUEUED commands are supported. There is considerable performance enhancement when performing I/O operations using the Marvell driver with Sun branded Hitachi model HDS7225SBSUN250G.

x86: NVIDIA Accelerated Graphics Drivers

This driver enhancement is new in the Developer 2/07 release.

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.

Adaptec aac Hardware Support

This driver enhancement is new in the Developer 2/07 release.

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: Direct Rendering Infrastructure Porting

This driver enhancement is new in the Developer 2/07 release.

In this release, the Direct Rendering Infrastructure (DRI) kernel-level framework and a kernel (Direct Rendering Manager) driver for Intel integrated chipsets have been ported from BSD to Solaris.

For more information about DRI, see “Direct Rendering Infrastructure” on page 80.

New Features in Solaris Express 12/06

This section describes all features that are new or have been enhanced in the Solaris Express 12/06 release.

Session Initiation Protocol Library (libsip)

This networking enhancement is new in the Solaris Express 12/06 release.

Session Initiation Protocol (SIP) is an application layer protocol that can be used to initiate, modify, and terminate multimedia sessions such as Voice-over-IP (VoIP) and Instant Messaging (IM).

SIP library in Solaris provides a SIP stack, that conforms to RFC 3261, and a set of API to write SIP applications. Primary users of this library include developers who write SIP applications such as soft phones, proxy servers, redirect servers, etc.

The library supports all SIP headers in RFCs 3261, 3262, 3265, 3323, and 3325.

For more information about SIP, see the sip(7P) man page.

ZFS Command History (zpool history)

This system administration tools enhancement is new in the Solaris Express 12/06 release.

ZFS automatically logs successful zfs and zpool commands that modify pool state information. For example:

zpool history

History for 'newpool':

```
2006-10-23.08:58:22 zpool create -f newpool c1t2d0
2006-10-23.08:59:02 zpool replace -f newpool c1t2d0 c1t3d0
2006-10-23.08:59:54 zpool attach -f newpool c1t3d0 c1t4d0
```

This feature enables you or Sun support personnel to identify the *exact* set of ZFS commands that was executed to troubleshoot an error scenario.

The features of the history log are as follows:

- The log cannot be disabled.
- The log is saved persistently on disk, which means the log is saved across system reboots.
- The log is implemented as a ring buffer. The minimum size is 128 Kbytes. The maximum size is 32 Mbytes.
- For smaller pools, the maximum size is capped at 1% of the pool size, where *size* is determined at pool creation time.
- The log does not require any administration. This means that you do not need to tune the size of the log or change the location of the log.

Currently, the `zpool history` command does not record *user-ID*, *hostname*, or *zone-name*.

For more information about troubleshooting ZFS problems, see *Solaris ZFS Administration Guide*.

Changes and Improvements to Removable Media Management

This device management enhancement is new in the Solaris Express 12/06 release.

Starting with this release, better services and methods that provide removable media management have replaced previous features for managing removable media.

The following new features are available:

- New removable media services are enabled and disabled by using SMF.

```
online      12:17:54  svc:/system/hal:default
online      12:17:56  svc:/system/filesystem/rmvolmgr:default
online      12:17:26  svc:/system/dbus:default
```

- Removable media is now mounted automatically in the `/media` directory. However, symbolic links to `/media` are provided from previous media mount points, `/cdrom` and `/rmdisk`, for compatibility purposes.

For example, a compact flash memory card (`/dev/dsk/c4d0p0:1`) is mounted as follows:

```
$ ls /media/NIKON
```

For example, a USB memory stick (`/dev/dsk/c3t0d0s0`) is mounted as follows:

```
$ ls /media/U3
```

For example, a diskette (`/dev/diskette0`) is mounted as follows:

```
$ ls /media/floppy
```

- The default removable media volume manager, `rmvolmgr`, is responsible for following activities:
 - Mounting and unmounting volumes.
 - The root instance of `rmvolmgr` starts at system boot. However, you can configure your session's configuration files to start an instance of `rmvolmgr` when you log in. When run within a user session, `rmvolmgr` only mounts devices owned by the current user or session and does not conflict with the root instance.
 - When `rmvolmgr` exits, it unmounts all media that it mounted.
 - For compatibility purposes, `rmvolmgr` creates symbolic links under the `/cdrom`, `/floppy`, `/rmdisk` directories to the actual mount points under `/media`.
 - A special `rmvolmgr` run mode is available for CDE compatibility.
- The hardware abstraction layer (HAL) daemon, `halld`, provides a view of the device attached to a system. This view is updated automatically as hardware configuration changes, by hot-plugging or other mechanisms.

HAL represents a piece of hardware as a device object. A device object is identified by a unique device identifier (UDI) and carries a set of key-value pairs referred to as device properties. Some properties are derived from the actual hardware, some are merged from device information files (`.fdi` files) and some are related to the actual device configuration.

The following features are removed:

- The `vol` daemon, the `vol` file system, and the `vol` service have been removed.

```
svc:/system/filesystem/vol
```

- Logical device names for removable media under the `/vol` directory, such as `/vol/dev/rdisk/...` or `/vol/dev/aliases/...`, are no longer provided.

To access removable media by its logical device name, the `/dev` device should be used. For example:

```
/dev/rdisk/c0t6d0s2
```

- Some `vol` device nicknames are no longer available. The following `ejct -l` output identifies the available device nicknames for each device and in the example, the mounted media pathnames (`/media/SOL_11_X86_4`):

```
$ eject -l
/dev/dsk/c2t0d0s2    cdrom,cdrom0,cd,cd0,sr,sr0,SOL_11_X86_4,/media/SOL_11_X86_4
/dev/diskette        floppy,floppy0,fd,fd0,diskette,diskette0,rdiskette,rdiskette0
```

The comma-separated list shows the nicknames that can be used to eject each device.

- Customizations that were made in `vol.d.conf` and `rmmount.conf` are no longer available because these configuration files no longer exist. For information about managing media customizations, see [“Customizing Removable Media Management” on page 85](#).
- Commands that begin with `vol*` commands except for `volcheck` and `volrmmount`.

Backward Compatibility

The following features provide backward compatibility with previous Solaris removable media features:

- Removable media mount points have moved to the `/media` directory, which is used to mount removable media, such as CD-ROMs and USB devices. Symbolic links to `/media` from previous media mounts points, such as `/cdrom` and `/rmdisk`, are provided for compatibility purposes.
- The `rmformat` command is still available. The output of this command is identical to what it looks in previous Solaris releases with `vol.d` disabled.

For example:

```
# rmformat
Looking for devices...
  1. Logical Node: /dev/rdisk/c0t6d0s2
     Physical Node: /pci@1f,4000/scsi@3/sd@6,0
     Connected Device: TOSHIBA DVD-ROM SD-M1401 1009
     Device Type: DVD Reader
     Bus: SCSI
     Size: 2.9 GB
     Label: <None>
     Access permissions: <Unknown>
```

- The `eject` command is available but has been enhanced. For more information, see [“Ejecting Removable Media” on page 85](#).

Mounting and Unmounting Removable Media

Most commands that begin with `vol*` are removed in this release. A modified version of `rmmount` and a new `rmumount` command are available to mount and unmount removable media.

These commands can be used to mount by device name, label, or mount point. For example, to mount an iPod:

```
% rmmount ipod
```

For example, to unmount the file systems on a DVD:

```
# rmount cdrom
cdrom /dev/dsk/c0t6d0s5 unmounted
cdrom /dev/dsk/c0t6d0s0 unmounted
```

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

Mounting and Unmounting Diskettes

You can use the existing `volcheck` command to manually poll diskettes and mount them if a new diskette is detected.

If you manually reformat diskette after it is connected to the system, HAL is not automatically notified. Continue to use the `volcheck` command to notify the system and attempt to automount a new file system on a diskette.

Ejecting Removable Media

As in previous Solaris releases, use the `eject` command to unmount and eject removable media. However, the following `eject` options are available:

- f Forces the device to eject even if the device is busy.
- l Displays paths and nicknames of devices that can be ejected.
- t A CD-ROM tray close command is provided to the device. Not all devices support this option.

For example, to eject by its volume label:

```
% eject mypictures
```

As in previous Solaris releases, you might need to issue the `volcheck` command before using the `eject` command to eject a diskette.

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

Customizing Removable Media Management

For most customizations that were available in the `vol.d.conf` and `rmount.conf` files, you will need to either use Desktop Volume manager preferences or modify the `.fdi` files.

- For `rmount.conf` actions, you will need to use either Desktop Volume Manager actions, `gconf`, or HAL callouts.
- Previously, `rmount.conf` actions could be run as root on behalf of ordinary users. Now, this is done by installing callout executables in the `/usr/lib/hal` directory.

Disabling Removable Media Features

You can disable some or all removable media features in this release:

- To prevent volumes from mounting outside of user sessions, disable the `rmvolmgr` service. For example:

```
# svcadm disable rmvolmgr
```

- To prevent any volume management, disable the `dbus`, `hal`, and `rmvolmgr` services.

```
# svcadm disable rmvolmgr
# svcadm disable dbus
# svcadm disable hal
```

Disabling these services means that you would have to mount all media manually by using the `mount` command.

SPARC: Process Count Scalability

This system resources enhancement is new in the Solaris Express 12/06 release.

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.

Packet Filter Hooks

This networking enhancement is new in the Solaris Express 12/06 release.

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.

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. Previous Solaris releases 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/>.

Name service switch enhancements

This release contains upward compatible changes to the name service switch (nss) and to the Name Switch Cache Daemon - nscd(1M)- in order to deliver new functionality including 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 Microsoft Active Directory.
- A framework for the future addition of putXbyY interfaces.

New Features in Solaris Express 11/06

This section describes all features that are new or have been enhanced in the Solaris Express 11/06 release.

Resource Controls to Limit Locked Physical Memory

This system resource enhancement is new in the Solaris Express 11/06 release.

To limit the amount of locked physical memory available to a zone on a Solaris OS with zones installed, use the `zone.max-locked-memory` zone-wide resource control. The resource control is set through the `add rctl` resource property in `zonecfg` for non-global zones. With the introduction of `zone.max-locked-memory`, the `proc_lock_memory` privilege is now part of the standard default set of zone privileges.

The allocation of the locked physical memory resource across projects within the zone can be controlled by using the `project.max-locked-memory` resource control.

The `project.max-locked-memory` resource control replaces the `project.max-device-locked-memory` resource control, which has been removed from the Solaris OS.

For more information, see the following:

- *System Administration Guide: Virtualization Using the Solaris Operating System*
- `zonecfg(1M)` man page
- `resource_controls(5)` man page

x86: lx Branded Zones: Solaris Containers for Linux Applications

This system resource enhancement is new in the Solaris Express 11/06 release.

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 `lx` 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 `lx` 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: Virtualization Using the Solaris Operating System*.

Also see the following man pages:

- `zoneadm(1M)`
- `zoncfg(1M)`
- `brands(5)`
- `lx(5)`

Name Service Switch and `nscd` Enhancements

This system resource enhancement is new in the Solaris Express 11/06 release.

Starting with this release, the name service switch is updated. The name service switch includes `nscd` and all `getXbyY` interfaces that query files and network data from DNS, NIS, NIS+, or LDAP. The behavior of the name service switch enhancements is identical to previous Solaris Express releases.

The major functional change in the Solaris Express 11/06 release is that, when you enable `nscd`, `nscd` performs all name service lookups. Prior to this release, `nscd` cached only a small subset of lookups. To enable `nscd`, type the following command:

```
# svcadm enable name-service-cache
```

Note – `nscd` is normally enabled by default.

If incorrect name service behavior such as incorrect `getXbyY` results or `nscd` hangs, is detected while `nscd` is running, either restarting or disabling `nscd` should correct the behavior. To restart `nscd` type the following command:

```
# svcadm restart name-service-cache
```

Disabling `nscd` automatically forces applications to perform all their own name service lookups as in prior releases of Solaris Express. To disable `nscd`, type the following command:

```
# svcadm disable name-service-cache
```

`nsd` will use a naming service, such as NIS, NIS+, or LDAP, only if the Service Management Facility (SMF) has enabled that service.

Forcibly Unmount the PCFS File System

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

Starting with this release, you can use the `-f` option with the `umount` command to forcibly unmount the PCFS file system.

New Features in Solaris Express 10/06

This section describes all features that are new or have been enhanced in the Solaris Express 10/06 release.

System V Resource Controls for Zones

This system resource enhancement is new in the Solaris Express 10/06 release.

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:

- `prctl(1)` man page
- `zonecfg(1M)` man page
- `resource_controls(5)`
- *System Administration Guide: Virtualization Using the Solaris Operating System*

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 enhancement was introduced in the Solaris Express 5/06 release.

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*

Single Hosts File

This networking enhancement is new in the Solaris Express 10/06 release.

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)` man page and the `ipnodes(4)` man page.

SPARC: New `sysidkbd` Tool Configures Your Keyboard

This installation enhancement is new in the Solaris Express 10/06 release.

For SPARC based platforms, a new `sysidtool`, `sysidkbd`, configures your USB keyboard layout during system installation.

Note – Previously, the USB keyboard attached on the SPARC platform always 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 during installation.

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

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

If the keyboard is not self-identifying and you want to prevent being prompted during your JumpStart installation, select the keyboard language in your `sysidcfg` script. For JumpStart installation, the default is for a U.S. English keyboard layout. To select another language, set the keyboard entry in your `sysidcfg` script as in this example:

```
keyboard=German
```

Note – The value provided for `sysidcfg` must be a valid value. Otherwise, an interactive response is required during installation. Valid keyboard strings are defined in a file which is referenced in the `sysidcfg(4)` man page.

For more information, see the `sysidcfg(4)` man page and the `sysidtool(1M)` man page.

Device Naming Enhancements

This device management enhancement is new in the Solaris Express 10/06 release.

Starting with this release, the `/dev` name space supports multiple file system instances as needed. When the system is booted, a global instance of the `/dev` file system is created automatically. Subsequent `/dev` instances are created and mounted when needed. For example, when devices are added to a non-global zone. When a non-global zone is shut down, the available `/dev` instance is unmounted and unavailable.

In addition, device configuration is improved in the following ways:

- **Reconfiguration boot is eliminated.** In release prior to Solaris Express 10/06, a reconfiguration boot was needed if you connected a device to a system that is powered off. Starting with this release, you do not need to perform a reconfiguration boot when attaching devices to a system that is powered off. When you reboot the system, the system automatically recognizes newly attached devices and creates the appropriate links.

For more information, see the `devfs(7FS)` man page.

- **Zone device support is simplified.** As described above, device support for Solaris zones is enhanced by providing specific instances of the `/dev` directory for non-global zones. In addition, zones are no longer dependent on the `devfsadm` daemon for reconfiguration of devices within a zone.
- **Pseudo device creation is improved.** Starting with this release, the content of the `/dev/pts` directory is created on demand in the global `/dev` name space and in the `/dev` instance when needed in a non-global zone. In addition, the `ptys` links are only visible in the global zone or the non-global zone from which they are allocated.

For more information, see the `grantpt(3C)` man page.

Linker and Libraries Updates

This developer tools enhancement is new in the Solaris Express 10/06 release.

Starting with this release, there is greater flexibility in executing an alternative link-editor with the `link-editor -z altexec64` option and the `LD_ALTEEXEC` environment variable.

Symbol definitions that are generated using `mapfiles` can now be associated to the executable and linking format (ELF) sections.

The link-editors now enable the creation of static thread-local storage (TLS) within shared objects. In addition, a backup TLS reservation is established to provide for limited use of static TLS within post-startup shared objects.

New Features in Solaris Express 9/06

This section describes all features that are new or have been enhanced in the Solaris Express 9/06 release.

MPSS Extension to Non-ISM/DISM Anonymous Shared Memory

This system performance enhancement is new in the Solaris Express 9/06 release.

Starting with this release, Solaris large-page support has been enhanced by extending multiple page size support (MPSS) to the following:

- Non-ISM/DISM SysV shared memory
- `MAP_SHARED` mappings created by `mmap()` of `/dev/zero` or with the `MAP_ANON` flag

In releases prior to Solaris Express 9/06, user applications could only map ISM/DISM SysV segments with large pages and there was no support for the use of large pages for segments created by `MAP_SHARED` `mmap()` of `/dev/zero` or `mmap()`. The new feature extends large-page support in the following two ways:

- On SPARC based systems, large pages are automatically assigned by the Solaris kernel to large enough shared memory mappings of either non-ISM/DISM SysV or `/dev/zero` or `MAP_ANON`.
- On SPARC and x86 based systems, the `mmap()` (`MC_HAT_ADVISE`) interface has been enhanced to enable users to explicitly request the use of large pages for `MAP_SHARED` anonymous memory (created by `mmap()` of `/dev/zero` or with the `MAP_ANON` flag) and non-ISM/DISM SysV memory.

In prior releases, however, `mmap()` (`MC_HAT_ADVISE`) against `MAP_SHARED` mappings worked only against memory created by mapping regular files. The significant advantage of the MPSS extension feature is potential performance improvement for applications that create large non-ISM/DISM shared memory segments or have big `MAP_SHARED /dev/zero` or `MAP_ANON` mappings. This performance improvement is due to the reduction of TLB misses due to the use of larger pages.

GNOME-VFS and Nautilus ACL Support

This desktop tools enhancement is new in the Solaris Express 9/06 release.

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.

ZFS Option to Sort List Output

This file system enhancement is new in the Solaris Express 9/06 release.

Starting with this release, the `zfs list` command has two new options, `-s` and `-S`. These options are used to select the columns that are to be displayed and sorted.

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

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.

Non-Global Zones Now Installed With Limited Networking Configuration

A non-global zone is now installed with limited networking configuration (`generic_limited_net.xml`). This means, for example, that only the `ssh` login is enabled by default, and logins through `rlogin` and `telnet` must be added if needed.

The administrator can switch the zone to the open, traditional networking configuration (`generic_open.xml`) by using the `net services` command, or enable and disable specific services by using the Service Management Facility (SMF) commands.

For more information about network configuration types, see Chapter 15, “Managing Services (Tasks)” in *System Administration Guide: Basic Administration*

For more information about the procedure to switch a non-global zone to a different network service configuration see, Chapter 22, “Logging In to Non-Global Zones (Tasks)” in *System Administration Guide: Virtualization Using the Solaris Operating System*

Adobe Flash Player Plug-in 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.

New Features in Solaris Express 8/06

This section describes all features that are new or have been enhanced in the Solaris Express 8/06 release.

OpenSSL 0.9.8a

This security enhancement is new in the Solaris Express 8/06 release.

OpenSSL 0.9.7d is upgraded to the stable OpenSSL 0.9.8a version. Some of the major changes in this version include:

- DTLS support
- Implementation of SHA-224/-256/-384/-512
- Re-implementation of big numbers (BIGNUM) support
- New STORE type as a common interface to certificates and key stores

- IPv6 support for certificate extensions

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

This language support enhancement is new in the Solaris Express 8/06 release.

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://unicode.org/cldr>.

hostname Command

This system administration tool enhancement is new in the Solaris Express 8/06 release.

Starting with this release, the `hostname` command has been re-implemented as a stand-alone binary executable program. The `hostname` command has also been enhanced with `getopt(3C)` support to detect and reject invalid command-line options.

Zone Unique Identifier

This system administration tool enhancement is new in the Solaris Express 8/06 release.

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”

This system administration tool enhancement is new in the Solaris Express 8/06 release.

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.

Changes to How \$TERM Value for Console is Set

This system administration enhancement is new in the Solaris Express 8/06 release.

The `$TERM` value is dynamically derived and depends on the terminal emulator that the console is using.

- On SPARC based systems the `$TERM` value is as follows:

<code>sun - color</code>	If the system uses the kernel's terminal emulator
<code>sun</code>	If the system uses the PROM's terminal emulator
- On x86 based systems, the `$TERM` value is `sun - color` because the kernel's terminal emulator is always used.

For more information, see the following:

- [“Coherent Console” on page 127](#) in the Solaris Express 3/06 release
- “Managing Terminals and Modems” in the *System Administration Guide: Advanced Administration*

Note – This change does not impact how the terminal type is set for the serial port. You can still use the `svccfg` command to modify the `$TERM` value.

Solaris Zones Boot Enhancements

This system resource enhancement is new in the Solaris Express 8/06 release.

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 `zoneadm(1M)` and `zonecfg(1M)` man pages
- *System Administration Guide: Virtualization Using the Solaris Operating System*

x86: Solaris Audio Driver for Ultra 20 M2 Workstation

This driver is new in the Solaris Express 8/06 release.

The `audiohd` driver is a Solaris High Definition audio driver for the Ultra 20 M2 Workstation. This new driver enables users to play “high definition” audio in the Ultra 20 M2 platform.

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

ld Link Editor for Object Files

This developer tool enhancement is new in the Solaris Express 8/06 release.

Starting with this release, the behavior of the `ld` command when creating a new version of an existing output file has been changed. This change addresses a long-term issue with the way new linker output files interact with running programs that are simultaneously using the old versions of those files.

If the file being created by `ld` already exists, the existing file is unlinked after all input files have been processed. A new file with the specified name is then created. This behavior enables the `ld` command to create a new version of the file, while simultaneously allowing existing processes that are accessing the old file contents to continue running. If the old file has no other links, the disk space of the file is freed when the last process referencing the file terminates.

In releases prior to Solaris Express 8/06, the existing file was not unlinked, and was instead over-written. However, this approach had the potential to corrupt any running processes that were using the file.

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

Note – The new behavior of the `ld` command has an implication for output files that have multiple hard links in the file system. Previously, all links would remain intact, with all links accessing the new file contents. The new `ld` behavior “breaks” such links, with the result that only the specified output file name references the new file. All the other links continue to reference the old file. To ensure consistent behavior on any Solaris version, applications that rely on multiple hard links to linker output files should explicitly remove and re-link the other file names.

Support for iSCSI Target Devices

This file system enhancement is new in the Solaris Express 8/06 release.

This Solaris release provides support for iSCSI target devices, which can be disk or tape devices. Releases prior to Solaris Express 8/06 provided support for iSCSI initiators. The advantage of setting up Solaris iSCSI targets is that, existing fibre channel devices 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 `iscsiadm` 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.

For more information, see:

- `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*

zfs snapshot **Command**

This file system enhancement is new in the Solaris Express 8/06 release.

Starting with this release, a new flag, `-r`, has been added to the `zfs snapshot` command. The new flag enables system administrators to take many snapshots at once. Using the `-r` flag, is easier than running multiple `zfs snapshot` commands and execution time is lessened.

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

Solaris iSCSI Initiator

This device management enhancement is new in the Solaris Express 8/06 release.

Starting with this release, the Solaris iSCSI Initiator has been extended to support the creation of multiple iSCSI sessions to different IP addresses in the same iSCSI target portal group. This feature extends the iSCSI initiator's MPxIO support for a large number of additional storage arrays.

In releases prior to Solaris Express 8/06, the Solaris iSCSI Initiator supported the following actions:

- Creation of multiple iSCSI sessions to many IP addresses in different iSCSI target portal groups
- Creation of multiple iSCSI sessions to the same IP address in the same iSCSI target portal group
- Creation of multiple iSCSI sessions to arrays supporting iSCSI login redirection

For more information see, Chapter 14, “Configuring Solaris iSCSI Targets and Initiators (Tasks),” in *System Administration Guide: Devices and File Systems*.

PAPI Print Commands

This desktop tools enhancement is new in the Solaris Express 8/06 release.

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:

- “Implementation of PAPI Print Client Commands” in *System Administration Guide: Solaris Printing*
- OpenSolaris Printing Community web pages at http://opensolaris.org/os/community/printing/projects/papi_client_commands/

Enhancements to `fstyp` Command

The following enhancements to the `fstyp` command are new in the Solaris Express 8/06 release.

Starting with this release, the `fstyp` command has a new option, `-a`, which displays the file system attributes in a consistent, name-value pair format. This command also supports DOS logical drive numbers. For example:

```
# fstyp /dev/dsk/c0t0d0p0:1
```

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

The following `fstyp` enhancements are also available starting with this release:

- Most of the `fstyp` functionality is now available to applications as a library API. For more information, see the `libfstyp(3LIB)` man page and the `3FSTYP` man page section.

- The `fstyp` script is now a binary that calls into `libfstyp(3LIB)` man page and the `3FSTYP` man page section. Existing `/usr/lib/fs/*/fstyp` back ends have been converted to `libfstyp(3LIB)` man page and the `3FSTYP` modules `/usr/lib/fs/*/fstyp.so.1`. The `/usr/lib/fs/*/fstyp` back ends are still available as links to `/usr/sbin/fstyp`.
- For backward compatibility, `fstyp` calls a legacy back end if an `fstyp.so.1` module is not found. Vendors are encouraged to convert their back ends to the new interface. For more information, see `fstyp_mod_init(3FSTYP)`.

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.

Starting with the Solaris Express 8/06 release, the Sun Java Web Console includes the following changes:

- 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)`
- 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 “Java Web Console” in *System Administration Guide: Basic Administration*.

Note – Starting with the Solaris Express 4/06, the Solaris ZFS web-based management tool is available in the Sun Java Web Console. This tool enables you to perform most of the administration tasks that you can perform with the command-line interface (CLI).

For more information about using the Solaris ZFS web-based management tool, see the *Solaris ZFS Administration Guide*.

New Features in Solaris Express 7/06

This section describes all features that are new or have been enhanced in the Solaris Express 7/06 release.

Ability to Validate a Non-Global Zone Migration Before the Migration Is Performed

This system resources enhancement is new in the Solaris Express 7/06 release.

Starting with this release, a zone migration dry run is performed before the actual non-global zone migration. The `zoneadm detach` subcommand can now generate a `manifest` on a running zone without actually detaching the zone. The `zoneadm attach` subcommand can then 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 about procedures for migrating a non-global zone to a different machine and performing a dry run before the actual move, see the following:

- `zoneadm(1M)` man page
- “Migrating A Non-Global Zone to a Different Machine” in the *System Administration Guide: Virtualization Using the Solaris Operating System*

Mailbox Size Enhancement

This networking enhancement is new in the Solaris Express 7/06 release.

The `mail.local` program delivers mail on the Solaris OS. In releases prior to Solaris Express 7/06, the mail client could support a maximum mailbox size of 2,147,483,647 bytes (or 2Gbytes - 1). This limitation has been removed. Now the mailbox size can be as large as any other Solaris supported file system.

For more information about `mail.local`, see the `mail.local(1M)` man page.

Exclusive Link-Based Failure Detection for IPMP Singleton

This networking enhancement is new in the Solaris Express 7/06 release.

Some sites use single-interface IP Multipathing (IPMP) groups in environments that cannot support probe-based network failure detection. Starting with this release, these sites can successfully deploy Solaris IPMP.

For more information about IPMP, see *System Administration Guide: IP Services*.

Runtime Linker Configuration File

This developer tool enhancement is new in the Solaris Express 7/06 release.

Runtime linker configuration files are created and managed with the `crle` command. These configuration files are used to alter default options for the Solaris runtime linker. In releases prior to Solaris Express 7/06, the runtime linker configuration files would encounter problems when used on AMD64 platforms. Starting with this release, the format of the runtime linker configuration files has been improved to resolve these problems. This improved format enables better file identification and ensures that the runtime linker does not use a configuration file generated on an incompatible platform.

Linker configuration files contain platform-specific binary data. A given configuration file can be interpreted by software with the same machine class and byte ordering. However, in releases prior to Solaris Express 7/06, the information necessary to enforce this restriction was not included in the linker configuration files.

Starting with this release, linker configuration files contain system identification information at the beginning of the file. This additional information is used by the `crle` command and the runtime linker to check the compatibility with linking configuration files. This information also allows the `file` command to properly identify linking configuration files. For backward compatibility, older linker configuration files will still be accepted but without the identification and error checks that are now available. When the update (`-u`) option is used on an older linker configuration file that lacks the system information, the `crle` command does not add system information to the result.

For more information, see the following man pages:

- `ld.so.1(1)`
- `crle(1)`

- `file(1)`

Solaris Fibre Channel Host-Based Logical Unit Number Masking

This device management enhancement is new in the Solaris Express 7/06 release.

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.

Solaris Trusted Extensions

This security enhancement is new in the Solaris Express 7/06 release.

Starting with this release, 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, see the README and html files in the `ExtraValue/Cobundled/tx` directory.

Network Services Startup

This security enhancement is new in the Solaris Express 7/06 release.

Starting with this release, the generic installation has been changed. This change ensures that all network services, except `ssh`, are either disabled or restricted to respond to local requests only. The change to the generic installation also minimizes potential vulnerabilities that might be targeted by remote attackers.

Additionally, the network services startup feature enables customers to use only those services that they require. All of the affected services are controlled by the Service Management Framework (SMF). Any individual service can be enabled using the `sycadm` and `syccfg` commands. The `net services` command can be used to switch the service startup behavior.

Gnome 2.14

This desktop tools enhancement is new in the Solaris Express 7/06 release.

Starting with this release, the Java DS includes the new Gnome 2.14 desktop. The Gnome 2.14 desktop introduces new features and enhances the performance of many of the commonly used applications.

Evince PDF and PostScript Viewer

This desktop tools enhancement is new in the Solaris Express 7/06 release.

Starting with this release, the Evince PDF and PostScript Viewer has been added to the Java DS.

Firefox 1.5

This browser enhancement is new in the Solaris Express 7/06 release.

Starting with this release, Firefox 1.5 has been added to the Java DS and will now be the default browser.

Thunderbird 1.5

This email enhancement is new in the Solaris Express 7/06 Release.

Starting with this release, Thunderbird 1.5 has been added to the Java DS and will now be the default email client.

Note – The email client Evolution is still included in the Java DS.

Upgrading the Solaris OS When Non-Global Zones Are Installed

This installation enhancement is new in the Solaris Express 7/06 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 upgrade to the Solaris 7/06 release using, either the Solaris interactive installation program or custom JumpStart to upgrade. For details about using the Solaris interactive installation program, see the *Solaris Express Installation Guide: Solaris Live Upgrade and Upgrade Planning*. The process of upgrading with non-global zones installed has some limitations.

- A limited number of custom JumpStart keywords is supported.
For a list of supported custom JumpStart keywords, see the *Solaris Express Installation Guide: Custom JumpStart and Advanced Installations*.
- 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.
For more information, see *Installing With the Solaris Installation Program (Tasks)*, in *Solaris Express Installation Guide: Basic Installations*.
- 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 `lucreate` command, the `luupgrade` command cannot upgrade a boot environment that has non-global zones installed. In that case, the upgrade fails and an error message is displayed.

The ZFS File System

The following file system enhancements are new in the Solaris Express 7/06 release.

- **ZFS Double Parity RAID** - ZFS pools can be created using double parity RAID known as “raidz2”. Each raidz2 stripe within the pool can sustain up to two simultaneous failures without losing data.

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

- **ZFS Clone Promotion** - ZFS clones can be “promoted” to the active head of the original file system from which the snapshot was taken. The ZFS clone promotion feature enables the original file system and the snapshot to be deleted while preserving the clone as an active file system. In releases prior to Solaris Express 7/06, a snapshot could not be deleted while the clone still existed.

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

- **ZFS Hot-Spare Support** - ZFS supports the use of hot-spares within a pool. A device within a pool fails when it is removed from the system. The failed device will be automatically replaced by an available hot-spare. When the original failed device is usable again, the spare device is returned to the pool of available hot-spares.

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

For more information about the ZFS file system, see [“ZFS Command Improvements and Changes” on page 137](#).

New Features in Solaris Express 6/06

This section describes all features that are new or have been enhanced in the Solaris Express 6/06 release.

IPsec Policy Bypass Option While Running the `ping` Command

This networking enhancement is new in the Solaris Express 6/06 release.

In this release, the `-b` option has been added to the `ping` command. This option enables administrators to bypass global IPsec policy for a specified ping connection.

The IPsec policy bypass option enables isolating network issues from IPsec policy issues without changing or disabling global IPsec rules. Because, this option only affects the specified ping socket, encrypted and clear text pings can be run simultaneously. Only superuser or a user granted suitable privileges can use the IPsec policy bypass option.

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

Extended File Descriptor Limits For `stdio`

This system resource enhancement is new in the Solaris Express 6/06 release.

The Extended file descriptor limits for `stdio` feature removes the 256 FILE structure limitation on 32-bit Solaris processes. This feature provides both binary relief and programmatic interfaces.

For more information, see the following man pages:

- `extendedFILE(5)`
- `enable_extended_FILE_stdio(3C)`
- `fopen(3C)`
- `fdopen(3C)`
- `popen(3C)`
- `stdio(3C)`

Cloning Non-Global Zones With ZFS Zonpaths and Other ZFS Enhancements

This system resource enhancement is new in the Solaris Express 6/06 release.

The `zoneadm clone` command now automatically uses ZFS clone to clone a zone, when both the source and the target zonpaths reside on ZFS and are in the same pool.

The `zoneadm clone` command takes a ZFS snapshot of the source zonpath and sets up the target zonpath. The snapshot is named `SUNWzoneX`, where `X` is a unique ID which is used to distinguish between multiple snapshots. The destination zone's zonpath is used to name the ZFS clone. A software inventory is performed so that a snapshot used at a future time can be validated by the system.

Note that you can still specify that the ZFS zonpath be copied instead of ZFS cloned if desired.

To clone a source zone multiple times, a new parameter added to `zoneadm` command allows you to specify that an existing snapshot should be used. The system validates that the existing snapshot is usable on the target.

The zone installation process now has the capability to detect when a ZFS file system can be created for a zone. The process for uninstalling a zone can detect when a ZFS file system in a zone can be destroyed. These steps are then performed automatically by `zoneadm`.

MDI Enabling and Disabling of a Path

This system administration feature is new in the Solaris Express 6/06 release.

The MDI enabling and disabling functionality provides a path management mechanism to enable or disable a path for a target device. These new interfaces to MDI provide a user to enable

or disable the path from a pHCI to a target-address and a client service address (LUN address in case of `scsi_vhci`). This feature can be used by diagnostic applications to ensure that no user I/O is routed on the path on which diagnostic operations are being performed.

For more information, see the `mpathadm(1M)` man page. See also *Multipathing Administration Guide*.

SIGEV_THREAD Event Notification

This language support enhancement is new in the Solaris Express 6/06 release.

The SIGEV_THREAD Event Notification feature implements the POSIX-specified SIGEV_THREAD notification mechanism. The SIGEV_THREAD notification mechanism calls a user-specified function in the context of a separate thread when an event occurs.

This feature also extends the implementation of the existing SIGEV_PORT notification mechanism to message queues. This extension enables uniform application of all Solaris-supported notification mechanisms to all interfaces that accept a `sigevent` structure to request asynchronous notification. The following Solaris notification mechanisms are supported:

- SIGEV_NONE
- SIGEV_SIGNAL
- SIGEV_THREAD
- SIGEV_PORT

For more information about the asynchronous notification mechanisms, see the `signal.h(3HEAD)` man page.

x86: Uncacheable Memory Access Support

This system performance enhancement is new in the Solaris Express 6/06 release.

On x86 and AMD64 architectures, modern graphics controller cards, such as NVIDIA or AGP require a mechanism to access uncacheable memory to enhance performance. The uncacheable memory access enables these graphic controller cards to keep coherency on the system with performance gain.

The uncacheable memory access support feature provides this functionality by using the `ddi` interfaces.

The following memory type are supported:

- Uncacheable (UC)
- Write-combining (WC)

For more information, see the following man pages:

- `ddi_dma_mem_alloc(9F)`
- `devmap_devmem_setup(9F)`

Predictive Self-Healing for PCI Express on x64 Systems

This system administration feature is new in the Solaris Express 6/06 release.

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/>

Default Desktop Session in `dtlogin`

This desktop enhancement is new in the Solaris Express 6/06 release.

Now, when a user logs into the Solaris Desktop for the first time, Java Desktop System (JDS) is the default desktop environment instead of the Common Desktop Environment (CDE). JDS has also become the default environment for users who chose a desktop environment on an older Solaris release that is no longer present in the 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.

New Features in Solaris Express 5/06

This section describes all features that are new or have been enhanced in the Solaris Express 5/06 release.

Support for Descriptive Names for Metadevices and Hot-spare Pools

This system performance enhancement is new in the Solaris Express 5/06 release.

The rules for defining names for metadevices and hot-spare pools have been extended to allow the use of alphanumeric characters. Previously, names for metadevices were restricted to the form “dXXX”, while hot-spare pools had to use the form “hspYYY”.

For more information about the acceptable names for metadevices and hot-spare pools, see the *Solaris Volume Manager Administration Guide*.

File System Monitoring Tool (fsstat)

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.

The following `fsstat` example shows how to display all ZFS file system operations.

```
$ fsstat zfs
new name name attr attr lookup rddir read read write write
file remov chng get set ops ops ops bytes ops bytes
8.26K 240K 6.34K 4.03M 6.33K 12.6M 482K 204K 1.24G 13.4K 363M zfs
```

The following `fsstat` example shows how to display file system operations for the `/export/home` mount point.

```
$ fsstat /export/home
new name name attr attr lookup rddir read read write write
file remov chng get set ops ops ops bytes ops bytes
0 0 0 972 0 224 22 0 0 0 0 /export/home
```

The following `fsstat` example illustrates how to display file system operations for all file system types.

```
$ fsstat -F
new name name attr attr lookup rddir read read write write
file remov chng get set ops ops ops bytes ops bytes
```

```

1.71K 1.16K 37 371K 562 1.90M 1.04K 151K 142M 21.9K 55.3M ufs
0 0 0 1.60K 0 2.93K 344 1.30K 761K 0 0 proc
0 0 0 0 0 0 0 0 0 0 0 nfs
8.31K 240K 6.37K 4.04M 6.48K 12.6M 482K 204K 1.25G 13.7K 365M zfs
29 12 12 37.4K 8 22.5K 28 1.04K 2.34M 120 76.1K lofs
10.8K 4.72K 5.43K 54.6K 49 19.8K 28 173K 610M 513K 487M tmpfs
0 0 0 338 0 0 0 44 12.2K 0 0 mntfs
1 1 1 429 1 14.3K 115 1.34K 1002K 0 0 nfs3
12 6 9 150 0 442 30 91 3.32M 12 69.2K nfs4
1 0 1 20.8K 0 20.5K 0 0 0 0 0 autofs

```

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

useradd **Default Shell**

These security enhancements are new in the Solaris Express 5/06 release.

In this release, the following new options have been added to the `useradd -D` utility:

- `-s` option. Enables the user to change the default shell.
- `-k` option. Enables the user to change the default `skel` directory.
- `-b` option. Enables the user to change the base directory.

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

ZFS Pool Import **Destroyed Pools**

This file system enhancement is new in the Solaris Express 5/06 release.

In this release, the `zpool` command can now reimport previously destroyed pools that still have their data intact.

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

For more information about the ZFS file system, see [“ZFS Command Improvements and Changes” on page 137](#).

Configurable Privileges for **Non-Global Zones**

These system resources enhancements are new in the Solaris Express 5/06 release.

The `zonecfg` command can now 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.

Note the following:

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

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

- *System Administration Guide: Virtualization Using the Solaris Operating System*
- `zonecfg(1M)` man page

SO_TIMESTAMP Socket Option

This networking enhancement is new in the Solaris Express 5/06 release.

The `SO_TIMESTAMP` option enables or disables the reception of a timestamp with datagram. If the `SO_TIMESTAMP` option is enabled on a `SOCK_DGRAM` or `SOCK_RAW` socket, the `recvmsg(2)` call returns a timestamp in the native data format, corresponding to when the datagram was received.

For more information, see the following man pages:

- `setsockopt(3XNET)`
- `getsockopt(3XNET)`
- `recvmsg(3XNET)`

NFSv4 Domain Name Configurable During Installation

This system administration feature is new the Solaris Express 5/06 release.

The NFS version 4 domain can now be defined during the installation of the OS. To facilitate this new functionality, the `sysidnfs4` program runs during the installation process to determine whether an NFSv4 domain has been configured for the network. In previous Solaris 10 releases, the NFS domain name was defined during the first system reboot after installation.

The NFSv4 domain can now be defined as follows:

- If you are using the Solaris interactive installation program, you can choose the default, which automatically derives the NFSv4 domain name. Or, you can specify a different NFSv4 domain.
- If you are using the Solaris JumpStart™ program, a new keyword is available in the `sysidcfg` file. You can now assign a value for the NFSv4 domain by using the new keyword, `nfs4_domain`.

For information about the NFSv4 domain name configuration	<i>System Administration Guide: Network Services</i>
For information about Solaris interactive installations	<i>Solaris 10 11/06 Installation Guide: Basic Installations</i>
For information about Solaris network installations	<i>Solaris 10 11/06 Installation Guide: Network-Based Installations</i>
For information about Custom JumpStart installations	<i>Solaris 10 11/06 Installation Guide: Custom JumpStart and Advanced Installations</i>
For information about the <code>sysid</code> command tools	<code>sysidtool(1M)</code> and <code>sysidnfs4(1M)</code> man pages

Using DTrace in a Non-Global Zone

This system resource enhancement is new in the Solaris Express 5/06 release.

DTrace can now be used in a non-global zone when the `dt race_proc` and `dt race_user` privileges are assigned to the zone. DTrace providers and actions are limited in scope to the zone. With the `dt race_proc` privilege, `fasttrap` and `pid` providers can be used. With the `dt race_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 115](#) 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: Virtualization Using the Solaris Operating System*
- *Solaris Dynamic Tracing Guide*
- `zonecfg(1M)` man page
- `dt race(1M)` man page

64-bit SPARC: Extended Message Signaled Interrupt Support for Fire-based Platforms

This system resource enhancement is new in the Solaris Express 5/06 release.

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 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 Chapter 8, “Interrupt Handlers,” in *Writing Device Drivers*.

IPsec Kernel Module Error Logging

These system administration enhancements are new in the Solaris Express 5/06 release.

Starting with this release, all IPsec kernel module policy failures and other errors will be logged using the `ipsec_rl_strlog()` function. The `ipsec_rl_strlog()` function also has the ability to limit number of error messages sent to the system log. This ability prevents the system log from being overloaded.

The minimum interval between messages can be viewed or configured using the `ndd` command:

```
# ndd -get /dev/ip ipsec_policy_log_interval
```

The value returned is in milliseconds.

The `ipsec_policy_log_interval` now consolidates all IPsec-related error logging into a single function. This function also enables administrators to completely disable the error logging, as follows:

```
# ndd -set /dev/ip ipsec_policy_log_interval 0
```

Note – After rebooting the system, you need to disable the IPsec logging again.

iSCSI Logout Support

This system administration feature is new in the Solaris Express 5/06 release.

The iSCSI log out support feature allows a user to log out 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, the target logs out and cleans up all related resources. If the target is in use, 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 used 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 more information, see the `iscsiadm(1M)` man page. See also *System Administration Guide: Devices and File Systems*.

iSCSI MS/T Support

This system administration feature is new in the Solaris Express 5/06 release.

The iSCSI Multiple Session per 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 as follows:

- `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 administrators with iSCSI arrays that support login redirection.

For more information, see:

- `iscsiadm(1M)` man page
- *System Administration Guide: Devices and File Systems*
- Solaris iSCSI Multipathing Blueprint
- iSCSI RFC 3720 at <http://www.ietf.org/rfc/rfc3720.txt?number=3720>

iSNS Client Support for iSCSI

This device management feature is new in the Solaris Express 5/06 release.

The Internet Storage Name Service (iSNS) client feature adds a new discovery option to the Solaris OS iSCSI software initiator. This option enables a user to use the iSNS to handle Internet Protocol SAN (IP-SAN) device discovery. This is off by default and is not platform-specific. The iSNS client introduces several `iscsiadm` command additions, modifications, and driver changes to handle iSNS discovery.

Users 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 more information about the new and modified command-line options, see the `iscsiadm(1M)` man page. See also *System Administration Guide: Devices and File Systems*.

SNIA Multipath Management API support

This system administration feature is new in the Solaris Express 5/06 Release.

The Storage Networking Industry Association (SNIA) Multipath Management API (MP API) defines standard interfaces for multipath devices, associated path discovery, and path administration on a host. This feature provides Sun's implementation of the SNIA MP API library for the `scsi_vhci` driver-based multipathing solution.

The SNIA MA API consists of the following components:

- `/usr/sbin/mpathadm` CLI
- `/lib/libMPAPI.so` common library
- `/lib/libmpscsi_vhci.so` plug-in for the `scsi_vhci` driver

The following areas have been extended to support the SNIA MP API feature:

- MDI
- SCSI
- `libdevinfo`
- `scsi_vhci` IOCTL

With this feature, the administrators can use the standards-based path administration for `scsi_vhci` multipath devices.

For more information see the `mpathadm(1M)` and `libMPAPI(3LIB)` man pages. See also the *Solaris Fibre Channel Storage Configuration and Multipathing Support Guide*.

New Features in Solaris Express 4/06

This section describes all features that are new or have been enhanced in the Solaris Express 4/06 release.

Common Agent Container

This system administration feature is new in the Solaris Express 4/06 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 the 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 as follows:

- 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*.

Predictive Self-Healing for x64 Systems

This system administration feature is new in the Solaris Express 4/06 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 Express 4/06 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>

Java DTrace API

This system administration feature is new in the Solaris Express 4/06 release.

The Java Dynamic Tracing (DTrace) API is an interface to the native DTrace library. This interface provides support to develop visualization tools written in Java. The Java DTrace API comes installed in `/usr/share/lib/java/dtrace.jar`. This API allows multiple consumers of DTrace data to run simultaneously within a single Virtual Machine for the Java platform (JVM™ machine). Each consumer of the Java DTrace API listens for probe data and enables you to request consistent snapshots of aggregation data at any time.

For more information, see the Javadoc™ API at `/usr/share/lib/java/javadoc/dtrace/api/index.html`.

mkdtemp() and mkstemp() Library Functions

The library functions `mkdtemp()` and `mkstemp()` have been added to Solaris Express 4/06. The `mkdtemp()` function enables creation of uniquely named directories. The `mkstemp()` function enables creation of uniquely named files with a specific suffix.

For more information, see the `mkdtemp(3C)` and `mkstemp(3C)` man pages.

Zone Migration in Solaris Containers Technology

This system administration feature is new in the Solaris Express 4/06 release.

This feature enables migration of non-global zones from one machine to another machine. The `zonecfg` and `zoneadm` commands have been modified to enable migration of non-global zones from one system to another. The migration procedure detaches a halted zone from its current location and attaches the zone to a new system.

The global zone on the target system must be running the following:

- The same release of the OS as the original host
- The same versions of OS 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.

Because there are a number of ways to make the `zonpath` available on the new host, the actual movement of the `zonpath` from one system to another is a manual process that is performed by the global administrator.

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

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

- `zonecfg(1M)` man page
- `zoneadm(1M)` man page
- *System Administration Guide: Virtualization Using the Solaris Operating System*

Sun Java Web Console

In the Solaris Express 4/06 release, an enhancement has been made to the Sun Java Web Console feature that was introduced in the Solaris Express 10/04 release. See [“ZFS Web-Based Management” on page 138](#)

In this release, the Solaris ZFS web-based management tool is available in the Sun Java web console.

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

Support for PCI Express (PCIe)

This feature has been modified in the Solaris Express 4/06 release.

This Solaris release provides support for the PCI Express (PCIe) interconnect for both SPARC and x86 based systems.

PCIe is designed to connect peripheral devices to the following applications:

- Desktop
- Enterprise
- Mobile
- Communication
- Embedded

The PCIe interconnect is an industry-standard, high-performance, serial I/O bus. For details on PCIe technology, go to <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 this 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                ethernet/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 also 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 hotplug support are provided on your system. In addition, carefully review the instructions for physically inserting or removing adapters on your system, and review the semantics of device auto-configuration, if applicable.

PostgreSQL for the Solaris OS

This additional software is new in the Solaris Express 4/06 release.

PostgreSQL is a relational database system provided in the open-source community. More than 15 years of active development and a proven architecture have earned PostgreSQL a reputation for reliability, data integrity, and accuracy.

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

ZFS Persistent Offline

This system administration enhancement is new in the Solaris Express 4/06 release.

In this release, the `zpool offline` command offline a device persistently by default. You can use the `-t` option to offline a device temporarily.

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

For more information about the ZFS file system, see “ZFS Command Improvements and Changes” on page 137.

New Features in Solaris Express 3/06

This section describes all features that are new or have been enhanced in the Solaris Express 3/06 release.

32-bit: The `gnome-pilot` Utility

This desktop feature is new in the Solaris Express 3/06 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 Operating System (Solaris OS). The `gnome-pilot` feature supports synchronization with the PalmOS 4.x and the PalmOS 5.x.

IP_NEXTHOP Socket Option

This networking enhancement is new in the Solaris Express 3/06 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.

Basic Registration 1.1

This system administration feature is new in the Solaris Express 3/06 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 has been changed. The procedure for registering Solaris software has also 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 *System Administration Guide: Basic Administration*.

Sun Update Connection

This system administration feature is new in the Solaris Express 3/06 release.

SunSM Update Connection, System Edition product now 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>.

AES Counter Mode

This security enhancement is new in the Solaris Express 3/06 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.

x86: SATA HBA Framework Support

This driver is new in the Solaris Express 3/06 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)`

Coherent Console

This console subsystem feature is new in the Solaris Express 3/06 release.

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.

x86: Xorg X Server Version 6.9

This X11 windowing feature enhancement is new in the Solaris Express 3/06 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.

Resource Pools Service FMRI

This system resource enhancement is new in the Solaris Express 3/06 release.

Resource pools and dynamic resource pools have been integrated into the Solaris Service Management Framework Facility (SMF). The dynamic resource pools are now enabled separately from the resource pools service.

The Fault Management Resource Identifier (FMRI) for dynamic resource pools service is: `svc:/system/pools/dynamic`. The resource pools service FMRI is: `svc:/system/pools`. The enabling and disabling mechanisms through `pooladm` command are also still available.

Note that when a system with pools enabled is upgraded, if an `/etc/pooladm.conf` file exists, the pools service is enabled.

For more information, see the *System Administration Guide: Virtualization Using the Solaris Operating System*. See also the following man pages:

- `pooladm(1M)`
- `pool(1M)`
- `smf(5)`
- `libpool(3LIB)`

32-bit: RealPlayer for the Solaris OS

This additional software is new in the Solaris Express 3/06 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 Java Desktop System (JDS) users.

Zero-CountryCode Keyboard Layout Support

This language support enhancement is new in the Solaris Express 3/06 release.

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.

Using Compact Flash as ATA Disk

This system resource enhancement is new in Solaris Express 3/06 release.

Compact Flash (CF) can be used as an Advanced Technology Attachment (ATA) disk with a CF->ATA adapter. Using CF as an ATA disk provides a reliable bootable program and data storage device. This feature also improves the flexibility and reliability of the storage solution for your PC or application.

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

Zone Move and Clone Features in the Solaris Containers Technology

This system resources enhancement is new in Solaris Express 3/06 release.

Two new subcommands, `move` and `clone`, have been added to the `zoneadm` command. Zone move and clone features enable the following operations:

- 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 the `zoneadm(1M)` man page. See also the *System Administration Guide: Virtualization Using the Solaris Operating System*.

New Features in Solaris Express 2/06

This section describes all features that are new or have been enhanced in the Solaris Express 2/06 release.

Solaris Flash Archives

This installment enhancement is new in the Solaris Express 2/06 release.

This Solaris Flash enhancement enables a user to create an archive that includes large files. The `flarc` create command creates a Solaris Flash archive that can contain individual files of 4-Gbyte or more. The available archive utilities are:

- The portable archive interchange, `pax` utility. This utility can create an archive without size limitations on individual files. The `pax` utility is the default archiving utility.
- The `cpio` archive utility. This utility is invoked by using the `-L cpio` option. The `cpio` utility can archive individual files not larger than 2-Gbyte or 4-Gbyte. The file size limitation depends on the version of the `cpio` utility.

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 `-L cpio` option.

For more information, see the `pax(1)` and the `cpio(1)` man pages. See also the *Solaris Express Installation Guide: Solaris Flash Archives (Creation and Installation)*.

New `fallocate` Utility

This feature is new in the Solaris Express 2/06 release.

The `fallocate` utility enables contiguous allocation of capacity on a supported file system. The UNIX file system support for `fallocate` provides higher performance for programs that are performing sequential access to data.

For more information, see the `fallocate(2)` man page.

IIIMF and Language Engines

These language support enhancements are new in the Solaris Express 2/06 release.

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 JDS panel. Select **Add to Panel -> Utility -> InputMethod Switcher** to add the input method switcher to the JDS panel.
- **Utility for iiim-properties** - This feature supports various input method preferences. Use one of the following methods to launch 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` on 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.

32-bit: Palm Synchronization with USB Port

This desktop feature is new in the Solaris Express 2/06 release.

The palm synchronization feature enables Palm devices 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 `gpi lotd -control -applet(1)` man page.

The logadm Utility

This system administration feature is new in the Solaris Express 2/06 release.

The `logadm` utility provides the `-l` option to rotate the log file 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.

New Features in Solaris Express 1/06

This section describes all features that are new or have been enhanced in the Solaris Express 1/06 release.

ZFS Web-Based Management Feature

In the Solaris Express 1/06 release, an enhancement has been made to the ZFS feature that was introduced in the Solaris Express 12/05 release. For a description of this ZFS enhancement, see [“ZFS Web-Based Management” on page 138](#).

x86: PCI Express Support on x86 Systems

This device management feature is new in the Solaris Express 1/06 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 may 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 uses 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. And, review the semantics of device auto-configuration, if applicable.

SSL Proxy Module

This security enhancement is new in the Solaris Express 1/06 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. Applications, such as web servers, can be enabled 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 `ksstlcfg(1M)` man page.

TCP_INIT_CWND Socket Option

This networking enhancement is new in the Solaris Express 1/06 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.

x64: Watchdog Timeout Feature

This additional software is new in the Solaris Express 1/06 release.

The Watchdog mechanism detects a system hang. This feature is a timer that is continually reset by a user application as long as the operating system and user application are running. While the watchdog timer is operating in application mode, an additional alarm function, Alarm 3, is available to generate alarms in case of critical problems in the user application.

Object Migration and Interoperability Enhancements for `pktool`

This security enhancement is new in the Solaris Express 1/06 release.

The `pktool` command enables users to manage PKCS#11 objects. New subcommands have been added that can be used to move, display, and delete PKCS#11 objects. The subcommands can also show which PKCS#11 tokens are available. These operations help migrate cryptographic objects to or from the default Sun Software PKCS#11 softtoken or from other PKCS#11-compliant tokens.

For further information, see the `pktool(1)` man page.

Deimos Cryptographic Accelerator

This driver support is new in the Solaris Express 1/06 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.

HBA Drivers

These driver enhancements have been added to the Solaris Express 1/06 release.

The following HBA drivers have been added to the Solaris OS to support the full family of QLogic and Emulex HBA products which includes the following 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 choice of selection with a single fibre channel implementation. The supported HBAs are Solaris-ready certified and also 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://www.emulex.com/ts/docoem/sun/10k.htm>
- <http://qlogic.com>

Driver Support for New STK Tape Drive 10000 “Titanium”

This driver support is new in the Solaris Express 1/06 release.

This release adds the `st` driver, a driver for SCSI tape devices. This feature provides support for the new StorageTek™ tape drive T10000A “Titanium.”

For further information, see the `st(7D)` man page.

x86: Driver Support for AMD64 Platform

This driver support is new in the Solaris Express 1/06 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.

USB-to-Serial Driver for Prolific Adapters

This driver is new in the Solaris Express 1/06 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 `usbprl(7D)` man page.

USB-to-Serial Driver for Keyspan Adapters

This driver is new in the Solaris Express 1/06 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.

New Features in Solaris Express 12/05

This section describes all features that are new or have been enhanced in the Solaris Express 12/05 release.

The ZFS File System

This file system enhancement is new in the Solaris Express 12/05 release.

This Solaris Express release includes ZFS, a new 128-bit file system. ZFS provides simple administration, transactional semantics, end-to-end data integrity, and immense scalability. ZFS is not an incremental improvement to existing technology. Rather, ZFS is a fundamentally new approach to data management.

ZFS uses a pooled-storage model that completely eliminates the concept of volumes. Thus, 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, ZFS detects the damage and uses another copy to repair the damaged copy.

ZFS is Easy to Administer

For system administrators, the greatest improvement of ZFS over traditional file systems is the ease of administration.

ZFS takes a single command to set up a mirrored storage pool and file system. For example:

```
# zpool create home mirror c0t1d0 c1t2d0
```

The above command creates a mirrored storage pool named `home` and a single file system named `home`. The file system is mounted at `/home`.

With ZFS, you can use whole disks instead of partitions to create the storage pool.

Then, 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, 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 Command Improvements and Changes

The following section describes recent improvements and changes to the ZFS command interface in the Solaris Express release.

- **Clearing device errors** – You can use the `zpool clear` command to clear error counts 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.
- **Compact NFSv4 ACL format** – 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 3 ACL formats. Use the `ls -V` command to display compact and positional ACL formats and the `ls -v` command to display verbose ACL formats.
- **Double Parity RAID-Z (raidz2)** – A replicated RAID-Z configuration can now have either single- 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** – 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)** – 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.

- **Recovering destroyed pools** – The `zpool import -D` command enables you to recover pools that were previously destroyed with the `zpool destroy` command.
- **Temporarily take a device offline** – 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.
- **Upgrading ZFS Storage Pools** (`zpool upgrade`) – 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.
- **ZFS backup and restore commands are renamed** – The `zfs backup` and `zfs restore` commands are renamed to `zfs send` and `zfs receive` to more accurately describe their function. The function of these commands is to save and restore ZFS data stream representations.
- **ZFS and zones improvements** – 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: Virtualization Using the Solaris Operating System*.
- **ZFS is integrated with Fault Manager** – A ZFS diagnostic engine is included that is capable of diagnosing and reporting pool failures and device failures. 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*.

ZFS Web-Based Management

The Solaris Express 1/06 release includes the 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 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 it 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 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 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
```

ZFS and Solaris Zones

The Solaris Zones partitioning technology supports ZFS components, such as adding 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

In this release, the following Solaris installation tool support is provided:

- Custom Solaris 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 an install client:
 - `setup_install_server`
 - `add_install_server`
 - `add_install_client`

- Solaris 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. Therefore, your existing ZFS file systems are not copied into the boot environment (BE).
- 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.

New Solaris ACL Model

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 ZFS files.

The main features of the new ACL model are as follows:

- 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 `setfacl` and `getfacl` 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 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.

IPv6 for IP Filter

This security enhancement is new in the Solaris Express 12/05 release.

Solaris IP Filter has been enhanced to include IPv6 packet filtering. IPv6 packet filtering can filter based on the source IPv6 address, destination address, `poolsl` containing IPv6 addresses, and IPv6 extension headers.

The `-6` option has been added to the `ipf` command for use with IPv6. Use the `-6` option to load and flush IPv6 packet filtering rules. The `ipstat` command also has a `-6` option, which is used to display IPv6 statistics.

Although there is no change to the command line interface for the `ipmon` and `ippool` commands, these commands also support IPv6. The `ipmon` command has been enhanced to accommodate the logging of IPv6 packets. The `ippool` command supports the creation of IPv6 pools.

You can use the `ipf6.conf` file to create packet filtering rule sets for IPv6. By default, the `ipf6.conf` configuration file is included in the `/etc/ipf` directory. As with the other filtering configuration files, the `ipf6.conf` file loads automatically during the boot process when it is stored in the `/etc/ipf` directory.

For further information about Solaris IP, see the *System Administration Guide: IP Services*.

Support for Descriptive Names in Solaris Volume Manager

This system administration enhancement is new in the Solaris Express 12/05 release.

Descriptive names for both volumes and hot spare pools enable system administrators to name volumes and hot spare pools by using any name that follows the naming guidelines. Additionally, the `-D` option has been added to the `metostat` command. This option enables the user to list volumes and to list hot spare pools with descriptive names.

For more information, see the *Solaris Volume Manager Administration Guide*.

PKCS #11 v2.20 Support in the Solaris Cryptographic Framework

This security enhancement is new in the Solaris Express 12/05 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 `digest(1)` and `mac(1)` man pages.

UDP and TCP Performance Enhancement

This networking enhancement is new in the Solaris Express 12/05 release.

Performance of both the TCP protocol and the UDP protocol has been enhanced in this release. The enhancements result in lower latency and higher throughput for both transmit and receive performance. Network applications yield better performance due to system performance improvements. In particular, those applications which heavily transmit and receive UDP packets or utilize TCP loopback connections show greatest benefit.

For more information, see the `ip(7P)`, `tcp(7P)`, and `udp(7P)` man pages. See also the *Solaris Tunable Parameters Reference Manual*.

Zone Rename Feature in Solaris Containers Technology

This system resources enhancement is new in the Solaris Express 12/05 release.

The zone name is now an attribute that can be set through the `zonecfg` command. See the `zonecfg(1M)` man page for the procedure to change the name of a zone.

Only zones in the configured state or the installed state can be renamed. For information about zones states, see the `zones(5)` man page.

Advanced DDI Interrupt Framework

This device management enhancement is new in the Solaris Express 12/05 release.

The Solaris OS provides a new DDI interrupt framework for registering interrupts and for deregistering interrupts. Support for Message Signaled Interrupts (MSIs) is also provided. Management interfaces enable you to manipulate priorities and capabilities, to perform interrupt masking, and to obtain pending information.

For a list of the new interfaces, see “Interrupt Functions” in *Writing Device Drivers*. For a list of interrupt functions that are obsolete and should not be used, see the “B-3 Deprecated Interrupt Functions” table within this section.

For more information, see Chapter 8, “Interrupt Handlers,” in *Writing Device Drivers*. See also the individual man pages for the interfaces. All man pages for these interfaces are in the 9F man page section.

VLAN Support with xge 10Gb Ethernet Driver

This networking enhancement is new in the Solaris Express 12/05 release.

VLAN support with `xge` enables system administrators to configure Virtual LAN over a 10-Gbyte Ethernet. This enhancement is the first VLAN implementation with 10-Gbyte Ethernet driver in the Solaris OS. The `dladm` command can be used to configure VLAN with great flexibility.

For further information and VLAN configuration instructions, see the `dladm(1M)` man page.

Core Kerberos Mechanism Resync with MIT 1.4

This security enhancement is new in the Solaris Express 12/05 release.

The Kerberos_V5 GSS mechanism library has been synced with MIT Kerberos 1.4.0. This version includes fine-grained multithread support.

SPARC: Adobe Acrobat Reader 7.0.1

This desktop enhancement is new in the Solaris Express 12/05 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.

rge Driver

This driver support is new in the Solaris Express 12/05 release. The rge driver provides support for Realtek RTL8169S/8110S Gigabit Ethernet devices.

For further information, refer to the rge(7D) man page.

New UTF-8 Locales

This language support enhancement is new in the Solaris Express 12/05 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.

New Features in Solaris Express 11/05

This section describes all features that are new or have been enhanced in the Solaris Express 11/05 release.

Improvements to Volume Management (vol1d)

This volume management feature is new in the Solaris Express 11/05 release.

In this release, removable media management is improved. Previously, `vol` did not create device links for removable devices that contain no media. Now, device links are properly created for devices that contain no media, similar to the following:

```
lrwxrwxrwx 1 root root    28 Jun 13 13:09 /vol/dev/aliases/cdrom0
-> /vol/dev/rdsk/c2t2d0/nomedia
```

Now, you can use the `cdw` command and the `rmformat` command to list devices that have no media when `vol` is running.

You can revert back to the previous `vol` behavior by changing the following support `nomedia` entry in the `/etc/vold.conf` file as follows:

```
support media
```

For more information, see the `vold.conf(4)` man page.

In addition, `vol` 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.

For more information using these volume management improvements, see “What’s New in Removable Media?” in *System Administration Guide: Devices and File Systems*.

`vol` is Managed by the Service Management Facility (SMF)

This volume management feature is new in the Solaris Express 11/05 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*.

UFS Utility Enhancements

This file system enhancement is new in the Solaris Express 11/05 release.

The UFS file system check utility, `fsck`, has been enhanced to include features from the FreeBSD 4.9 version of the `fsck` program, as well as other enhancements.

The `fsck` utility in this Solaris release includes the following improvements:

- Checks and repairs file systems more thoroughly and provides improved error messages. For example, in some scenarios, `fsck` determines what structures are missing and replaces them appropriately.
- Automatically searches for backup superblocks.
- Reports when `fsck` needs to be rerun.
- When clearing directories, `fsck` now attempts to recover directory contents immediately and therefore, reduces the time spent rerunning this utility.
- If `fsck` finds duplicate blocks, and not all files that reference the duplicate blocks were cleared, `fsck` reports the inode numbers at the end of the `fsck` run. Then, you can use the `find` command to review the inodes that are damaged.
- Improved error messages regarding the status of extended attributes and other special files, such as device files and ACL entries, are included.
- Includes a `-v` option to enable more verbose messages.

In addition, the `newfs` and `mkfs` commands have been updated to include new options for displaying a file system's superblock information in text or dumping the superblock information in binary format.

```
newfs [ -S or -B ] /dev/rdisk/...
```

-S Displays the file system's superblock in text

-B Dumps the file system's superblock in binary

```
mkfs [ -o calcsb or -o calcbinsb ] /dev/rdisk/... size
```

-o calcsb Displays the file system's superblock in text

-o calcbinsb Dumps the file system's superblock in binary

The `fsck` utility uses this superblock information to search for backup superblocks.

For more information about these enhancements, see Chapter 17, “Managing File Systems (Overview),” in *System Administration Guide: Devices and File Systems*.

Internet Key Exchange (IKE) Enhancement

This security enhancement is new in the Solaris Express 11/05 release.

The `ikecert` utility now can copy keys and certificates that are on the disk to a hardware keystore device. The utility can also link existing key objects and certificate objects to the IKE database.

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

`cdrecord`, `readCD`, and `cdda2wav` Available

This device management feature is new in the Solaris Express 11/05 release.

Previously, `cdrecord` was available on 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` does. `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`, `cdda2wav` man pages in the `/usr/share/man` directory.

pilot-link Software

This Open Source software is new in the Solaris Express 11/05 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 <http://www.pilot-link.org>. See also `pilot-xfer(1)` in the `/usr/sfw/man` directory.

New Features in Solaris Express 10/05

This section describes all features that are new or have been enhanced in the Solaris Express 10/05 release.

mediaLib 2.3

This feature is new in the Solaris Express 10/05 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>.

New Features in Solaris Express 9/05

This section describes all features that are new or have been enhanced in the Solaris Express 9/05 release.

Dynamic Interrupt Distribution

This system performance enhancement is new in the Solaris Express 9/05 release.

The `int rd` daemon monitors loads on the system that are due to interrupts. If an imbalance is detected, the daemon retargets interrupts to new CPUs in order to balance the interrupt load. This feature should improve system performance under high interrupt loads.

The daemon is started by the new SMF service, `svc:/system/int rd`.

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

New Features in Solaris Express 8/05

This section describes all features that are new or have been enhanced in the Solaris Express 8/05 release.

Set SMTP to Use Transport Layer Security

This security enhancement is new in the Solaris Express 8/05 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*.

Driver for IEEE-1394 Based (IIDC) Digital Camera

This driver is new in the Solaris Express 8/05 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.

Refer to the `dcam1394(7D)` man page for more information.

SCSI HBA Driver for LSI MegaRAID 320-2x SCSI RAID Controller

This driver is new in the Solaris Express 8/05 release.

This release introduces the `lsimega` driver. The LSI 320-2x card is supported by `lsimega` with back-end support from LSI. 2. Sun Fire™ V20z/V40z servers with the optional 320-2x RAID card receive better support for RAID disk I/O and pass through I/O for SCSI tape and CDROM 3.

For further information, see the `lsimega(7D)` man page.

USB CCID IFD Handler

This driver is new in the Solaris Express 8/05 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.musclegcard.com>.

New Features in Solaris Express 7/05

This section describes all features that are new or have been enhanced in the Solaris Express 7/05 release.

x86: AGPgart Driver for x86 Systems

This driver is new in the Solaris Express 7/05 release.

AGPgart driver is an xserver kernel module that uses system memory to enhance graphics display. This feature is useful for memoryless graphics devices such as Intel integrated graphics and AGP graphics, which need to use system memory as video buffers.

Desktop users who use Intel 810/830/855 series graphics devices can achieve higher resolution of screen images under Xorg xserver with only 1 megabyte of BIOS-allocated video RAM. Video driver developers can utilize the AGPgart driver to allocate system memory for 2-D display or 3-D display and to manage AGP devices.

For more information, see the `agpgart_io(7I)` man page.

x86: New `prtconf` Option to Display Product Names

This system administration tool enhancement is new in the Solaris Express 7/05 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*.

New Features in Solaris Express 6/05

This section describes all features that are new or have been enhanced in the Solaris Express 6/05 release.

x86: GRUB-Based Booting

This system performance enhancement is new in the Solaris Express 6/05 release.

The Solaris OS is now loaded with the open source GRand Unified Bootloader (GRUB). GRUB is responsible for loading a boot archive, which contains the kernel modules and configuration files, into the system's memory. The Solaris kernel boots based on the contents of the in-memory boot archive.

Because the Solaris kernel is fully compliant with the Multiboot Specification, you can boot the Solaris OS on an x86 based system by using GRUB. With GRUB, you can more easily boot the various operating systems that are installed on your system. For example, on one x86 system, you could boot the following operating systems:

- Solaris OS
- Linux
- Microsoft Windows

A key benefit of GRUB is that it is intuitive about file systems and kernel executable formats. This feature enables you to load an operating system without recording the physical position of the kernel on the disk. With GRUB-based booting, the kernel is loaded by specifying its file name, the drive, and the partition where the kernel resides.

One notable change is the replacement of the Device Configuration Assistant (DCA) with the GRUB menu. When an x86 based system boots, the GRUB menu is displayed. From this menu, you can select an OS instance by using the up-arrow and down-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 on x86 based systems
- Installation from USB CD or DVD drives
- Ability to boot from USB storage devices
- Simplified DHCP setup for PXE boot with no vendor-specific options
- Elimination of all realmode drivers

In addition, the following two administrative commands have been added to assist with managing system bootability:

`bootadm` This command rebuilds the boot archive.
`installgrub` This command installs GRUB boot blocks.

Note – As a part of this functionality, x86 based systems must have a minimum of 256 Mbytes of RAM to boot and install the Solaris Express 6/05 release.

For more information about these enhancements, see *System Administration Guide: Devices and File Systems* and *System Administration Guide: Basic Administration*. See also the `bootadm(1M)`, `grub(5)`, and `installgrub(1M)` man pages.

See the following resources for additional installation information:

- For new booting information about installing with the Solaris interactive installation program, see the *Solaris Express Installation Guide: Basic Installations*.
- For new booting information that affects setting up an install server and installing over the network, see the *Solaris Express Installation Guide: Network-Based Installations*.
- For new booting information about installing with the custom JumpStart program, see the *Solaris Express Installation Guide: Custom JumpStart and Advanced Installations*.

Large Pages Enhancement

This system performance enhancement is new in the Solaris Express 6/05 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.

Large Pages for Kernel Memory

This system performance enhancement is new in the Solaris Express 6/05 release.

For 64-bit platforms, this feature supports mapping of the portion of the kernel heap with large pages. This feature increases performance of the system by reducing the number of Translation Lookaside Buffer (TLB) misses and the number of map operations or unmap operations.

Kernel Page Relocation

This system performance enhancement is new in the Solaris Express 6/05 release.

In this release, system performance has been enhanced on the Sun Fire 15K, Sun Fire 20K, and Sun Fire 25K domains.

Enhancements for bge and xge Network Interfaces

This networking enhancement is new in the Solaris Express 6/05 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.

Source-Filtered Multicasting

This networking enhancement is new in the Solaris Express 6/05 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*.

Automatic Renewal of Kerberos Credentials

This security enhancement is new in the Solaris Express 6/05 release.

A new option has been added to the `/etc/warn.conf` file. This option enables the `kttk_warn` daemon to automatically renew credentials for users who are logged in. If the `renew` option is

set, the users no longer need to run the `kinit -R` command in order to renew their credentials. In addition, options have been added to the `/etc/warn.conf` file that enable you to log the result when a renewal attempt is made.

See the `warn.conf(4)` man page for more information.

Enhancement to `ikecert certlocal` Command

This security enhancement is new in the Solaris Express 6/05 release.

The `ikecert certlocal` command now has options to set the validity period for certificate requests and for self-signed certificates. For more information, see the `ikecert(1M)` man page.

Enhanced Output With the `metaimport` Command

This system administration enhancement is new in the Solaris Express 6/05 release.

The Solaris Volume Manager `metaimport -r` command has been enhanced. The command shows the creation time for a disk set. The command also provides an advisory message when a disk is found in more than one disk set. These changes assist system administrators in determining which disk set to import.

For more information, see the *Solaris Volume Manager Administration Guide*.

Support for SCSI Disks Larger Than 2 Terabytes

This system administration enhancement is new in the Solaris Express 6/05 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*.

Advanced DDI Interrupts

This developer tool enhancement is new in the Solaris Express 6/05 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_cooki`
 - `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.

Door Interface Revisions

This developer tool enhancement is new in the Solaris Express 6/05 release.

Two new functions, `door_setparam` and `door_getparam`, have been added to the doors interfaces. These functions enable door servers to set boundaries on the data size and on the number of descriptors that are passed to `door_call` on the door. These enhancements simplify the door's handler function.

The `door_create` interface has a new flag, `DOOR_NO_CANCEL`. This flag disables the cancellation process of the server thread when the client aborts a `door_call` function due to a signal. For door servers which do not take advantage of cancellation, this change can prevent unwanted `EINTR` returns from interrupted system calls.

For more information, see the `door_setparam(3C)` and `door_create(3C)` man pages.

Memory Placement Optimization Hierarchical Lgroup Support

This performance enhancement is new 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 describes the API that can be used for the locality group (lgroup) abstraction. For more information, see `liblgrp(3LIB)`.

Support for Virtual USB Keyboard and Mouse Devices

This desktop enhancement is new in the Solaris Express 6/05 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)`.

X Client Support for XFree86 Extensions

These X11 windowing enhancements are new in the Solaris Express 6/05 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 |

`xvidtune` Provides video mode tuner for Xorg through XF86VidMode extension
`xvinfo` 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.

NVIDIA CK8-04 GE Driver Support

This driver support is new in the Solaris Express 6/05 release.

This release provides driver support for NVIDIA's gigabit Ethernet on the x86 platform. This features supports NVIDIA's Nfore4 chipset, CK8-04.

For further information, see the `nge(7D)` man page.

New Features in Solaris Express 4/05

This section describes all features that are new or have been enhanced in the Solaris Express 4/05 release.

TCP Keepalive Tunables

This networking enhancement is new in the Solaris Express 4/05 release.

With this enhancement, application developers can now fine tune the Transmission Control Protocol (TCP) keepalive mechanism on a per-socket basis. The `tcp(7P)` man page explains in detail how to use this enhancement.

New Features in Solaris Express 3/05

This section describes all features that are new or have been enhanced in the Solaris Express 3/05 release.

Additional Banner Page Printing Option in Solaris Print Manager

This system administration tool is new in the Solaris Express 3/05 release.

Solaris Print Manager has been expanded to include an additional “Never Print Banner” option. Using 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: Solaris Printing*.

New Features in Solaris Express 2/05

This section describes all features that are new or have been enhanced in the Solaris Express 2/05 release.

Support for iSCSI Devices

This device management feature is new 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.

Fibre-Channel HBA Port Utility

This system administration feature is new in the Solaris Express 2/05 release.

`fcinfo` is a command-line interface that collects administrative information on fibre-channel host bus adapter (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.

Metaslot in the Cryptographic Framework

This security feature is new 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 security enhancements are new 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 soft token 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 soft token keystore.

For further information about IKE, see the *System Administration Guide: IP Services*.

Xorg Release 6.8.2

This desktop enhancement is new 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.

New and Updated Drivers

In the Solaris Express 2/05 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.

Key Features in Solaris Express

The Solaris Express releases introduce enhancements to the Solaris 10 OS. The Solaris 10 OS includes the following key features.

In the Solaris 10 OS, Sun Microsystems 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. The Predictive Self-Healing feature includes significant changes to the booting and service administration processes. Also in the Solaris 10 OS, changes in the installation of the operating system provide a simplified and unified installation process.

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. Zones provide an isolated and secure environment for running applications.

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. 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. Also of particular importance, the networking stack for TCP connections has been rearchitected to deliver extremely high performance while improving the scalability.

Java 2 Platform, Standard Edition 5 is another key technology. Also of particular significance, the Solaris 10 software introduces support for the 64-bit computing capabilities of the AMD Opteron processor. Finally, the Solaris 10 software introduces the Java Desktop System. This desktop system combines open source software with Sun innovation.

For a complete summary of all the new features that are new in the Solaris 10 release, see *Solaris 10 What's New* at <http://docs.sun.com>. 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>.

