What's New in Solaris Express

Beta

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

Part No: 819–2183–22
December 2006

061207@15821

Contents

Preface

What's New in Solaris Express 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/`)
- Support (`http://www.sun.com/support/`)
- 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 `.login` file. Use `ls -a` to list all files. `machine_name% you have mail.`
`AaBbCc123`	What you type, contrasted with onscreen computer output	`machine_name%` **`su`** `Password:`
aabbcc123	Placeholder: replace with a real name or value	The command to remove a file is `rm` *filename*.

TABLE P–1 Typographic Conventions *(Continued)*

Typeface	Meaning	Example
AaBbCc123	Book titles, new terms, and terms to be emphasized	Read Chapter 6 in the *User's Guide*. A *cache* is a copy that is stored locally. Do *not* 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 12/06 release.

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, `hald`, provides a view of the device attached to a system. This view is updated automatically as hardware configuration changes, by hotplugging 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 `vold` daemon, the `volfs` file system, and the `volfs` service have been removed.

  ```
  svc:/system/filesystem/volfs
  ```

- Logical device names for removable media under the `/vol` directory, such as `/vol/dev/rdsk/...` 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/rdsk/c0t6d0s2
  ```

- Some `vold` device nicknames are no longer available. The following `eject -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 `vold.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 17.
- 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 `vold` disabled.

  For example:

```
# rmformat
Looking for devices...
     1. Logical Node: /dev/rdsk/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 17.

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:

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

For more information, see the `rmmount`(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 `vold.conf` and `rmmount.conf` files, you will need to either use Desktop Volume manager preferences or modify the `.fdi` files.

- For `rmmount.conf` actions, you will need to use either Desktop Volume Manager actions, `gconf`, or HAL callouts.
- Previously, `rmmount.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/`.

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: Solaris Containers-Resource Management and Solaris Zones*
- `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: Solaris Containers-Resource Management and Solaris Zones*.

Also see the following man pages:

- `zoneadm`(1M)
- `zonecfg`(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 the 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
```

`nscd` 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: Solaris Containers-Resource Management and Solaris Zones*

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: Advanced Administration*.

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: Advanced Administration*

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 `sysidkdb` Tool Prevents Having to Configure Your Keyboard

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

For SPARC based platforms, a new `sysidtool`, `sysidkdb`, 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 `sysidkdb` 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.

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 `memcntl(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, `memcntl(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 `netservices` 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: Solaris Containers-Resource Management and Solaris Zones*

Adobe Flash Player Plugin for Solaris

The Adobe Flash Player, formerly known as Macromedia Flash Player is the standard for delivering high-impact and rich web content. Designs, animation, and application user interfaces are deployed immediately across all browsers and platforms, attracting and engaging users with a rich web experience.

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://www.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:

  `sun-color` If the system uses the kernel's terminal emulator

  `sun` 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 58 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: Solaris Containers-Resource Management and Solaris Zones*

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

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

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:

- "PAPI Print Command and IPP Implementation" in *System Administration Guide: Advanced Administration*
- 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: Solaris Containers-Resource Management and Solaris Zones*

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 `netservices` 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 Sun Java™ Desktop System (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 68.

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 Zonepaths 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 zonepaths reside on ZFS and are in the same pool.

The `zoneadm clone` command takes a ZFS snapshot of the source zonepath and sets up the target zonepath. The snapshot is named SUNWzone*X*, where *X* is a unique ID which is used to distinguish between multiple snapshots. The destination zone's zonepath 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 zonepath 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 `defualtDt` 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 Hotspare Pools

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

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

For more information about the acceptable names for metadevices and hotspare 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 68.

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: Solaris Containers-Resource Management and Solaris Zones*
- `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	*System Administration Guide: Network Services*
For information about Solaris interactive installations	*Solaris 10 11/06 Installation Guide: Basic Installations*
For information about Solaris network installations	*Solaris 10 11/06 Installation Guide: Network-Based Installations*
For information about Custom JumpStart installations	*Solaris 10 11/06 Installation Guide: Custom JumpStart and Advanced Installations*
For information about the `sysid` command tools	`sysidtool`(1M) and `sysidnfs4`(1M) 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 `dtrace_proc` and `dtrace_user` privileges are assigned to the zone. DTrace providers and actions are limited in scope to the zone. With the `dtrace_proc` privilege, `fasttrap` and `pid` providers can be used. With the `dtrace_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 46 provides an overview of privileges in a non-global zone.

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

- *System Administration Guide: Solaris Containers-Resource Management and Solaris Zones*
- *Solaris Dynamic Tracing Guide*
- `zonecfg`(1M) man page
- `dtrace`(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-Experss 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 logout from an iSCSI target without rebooting the host. When a user tries to remove or disable a discovery method or address and the target is not in use, 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
- SCSA
- `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 `SUNCacaort` 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 21, "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 `mkstemps()` Library Functions

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

For more information, see the `mkdtemp`(3C) and `mkstemps`(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 zonepath available on the new host, the actual movement of the zonepath 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: Solaris Containers-Resource Management and Solaris Zones*

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 69

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 68.

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.

Sun[SM] Update Connection, System Edition product now supports the default patchadd 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 queueing 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 FMRIs

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: Solaris Containers-Resource Management and Solaris Zones*. See also the following man pages:

- `pooladm`(1M)
- `poold`(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: Solaris Containers-Resource Management and Solaris Zones*.

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 `flarcreate` 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 `gpilotd-control-applet`(1) man page.

The `logadm` Utility

This system administration feature in 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 69.

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 webservers, 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 l library for the others.

For more information, see the `ksslcfg`(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 `usbsprl`(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: Solaris Containers-Resource Management and Solaris Zones*.
- **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 `metstat` 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 (`vold`)

This volume management feature is new in the Solaris Express 11/05 release.

In this release, removable media management is improved. Previously, `vold` 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 `cdrw` command and the `rmformat` command to list devices that have no media when `vold` is running.

You can revert back to the previous `vold` 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, `vold` is now hot-plug aware. This improvement means that if you insert removable media, the media is automatically detected and mounted by `vold`. You do not need to restart `vold` 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*.

`vold` 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, `vold`, 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/rdsk/...*

`-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/rdsk/... 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 16, “Managing File Systems (Overview),” in *System Administration Guide: Devices and File Systems*.

Internet Key Exchange (IKE) Enhancment

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 intrd 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/intrd.

For more information, see the intrd(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.musclecard.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 `ktkt_warnd` 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:

`xgamma`	Alters a monitor's gamma correction through XF86VidMode extension
`xrandr`	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: Advanced Administration*.

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 `softtoken` keystore for applications that use the metaslot. When IKE uses the metaslot, you have the option of storing the keys on an attached board or in the `softtoken` keystore.

For further information about IKE, see the *System Administration Guide: IP Services*.

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`.