



Solaris 10 10/09 What's New



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

Part No: 821-0382-11
October 2009

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

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

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

This distribution may include materials developed by third parties.

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

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

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

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

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

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

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

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

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

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

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

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

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

Contents

Preface	5
1 What's New in the Solaris 10 10/09 Release	9
System Administration Enhancements	9
Two-Terabyte Disk Support for Installing and Booting the Solaris OS	9
pci tool Utility	9
ZFS Features and Changes	10
nss_ldap shadowAccount Support	13
Sun Validation Test Suite 7.0 Patch Set 6	13
Installation Enhancements	14
Turbo-Charging SVR4 Packaging	14
System Resources Enhancements	14
Zones Parallel Patching	14
PVIO Drivers in Solaris 10 Guest Domains	14
Device Management Enhancements	15
iSCSI Initiator SMF Service	15
LSI 6180 Controller Support in Solaris MPxIO	15
System Performance Enhancements	15
Callout Subsystem Scalability	15
Driver Enhancements	15
Enhanced 10 Gigabit Ethernet Performance	15
InfiniBand Feature Improvements	16
Driver for LSI MPT 2.0-Compliant SAS2.0 Controllers	17
x86: Broadcom NetXtreme II Gigabit Ethernet Driver	17
x86: Interrupt Remapping Support for Intel Vt-d	17
x86: SATA Tape Devices Supported by the AHCI Driver	17
Sun StorageTek 6Gb/s SAS PCIe RAID HBA Driver	18
Intel 82599 10Gb PCI Express Ethernet Controller	18

Intel 82598 10Gb PCI Express Ethernet Controller	18
Freeware Enhancements	18
NTP Version 4.2.5	18
PostgreSQL	18
Samba	19

Preface

Solaris 10 10/09 What's New summarizes all features in the Solaris™ 10 Operating System (OS) that are new or have been enhanced in the Solaris 10 10/09 OS.

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

In this document these x86 related terms mean the following:

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

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

Who Should Use This Book

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

Optional Feature Licensing

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

Related Books

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

Related Third-Party Web Site References

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

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

Documentation, Support, and Training

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

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

Sun Welcomes Your Comments

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

Typographic Conventions

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

TABLE P-1 Typographic Conventions

Typeface	Meaning	Example
AaBbCc123	The names of commands, files, and directories, and onscreen computer output	Edit your <code>.login</code> file. Use <code>ls -a</code> to list all files. <code>machine_name% you have mail.</code>
AaBbCc123	What you type, contrasted with onscreen computer output	<code>machine_name% su</code> Password:
<i>aabbcc123</i>	Placeholder: replace with a real name or value	The command to remove a file is <code>rm filename.</code>

TABLE P-1 Typographic Conventions (Continued)

Typeface	Meaning	Example
<i>AaBbCc123</i>	Book titles, new terms, and terms to be emphasized	Read Chapter 6 in the <i>User's Guide</i> . A <i>cache</i> is a copy that is stored locally. Do <i>not</i> save the file. Note: Some emphasized items appear bold online.

Shell Prompts in Command Examples

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

TABLE P-2 Shell Prompts

Shell	Prompt
C shell	machine_name%
C shell for superuser	machine_name#
Bourne shell and Korn shell	\$
Bourne shell and Korn shell for superuser	#

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

This document summarizes all features in the Solaris 10 Operating System (OS) that are new or have been enhanced in the current release, the Solaris 10 10/09 release.

For a summary of all features in the Solaris 10 OS that were introduced or enhanced since the Solaris 9 OS was originally distributed in May 2002, see [Solaris 10 What's New](#).

System Administration Enhancements

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

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

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

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

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

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

pcitool Utility

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

ZFS Features and Changes

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

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

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

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

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

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

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

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

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

```
# zpool create pool mirror c0t2d0 c0t4d0 cache c0t0d0
# zpool status pool
pool: pool
state: ONLINE
```

```
scrub: none requested
config:
```

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

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

For information about determining whether using cache devices is appropriate for your environment, see the *Solaris ZFS Administration Guide*.

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

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

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

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

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

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

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

- You can use the space usage properties to identify space usage for clones, file systems, and volumes, but not snapshots. The properties are as follows:
 - `usedbychildren` – Identifies the amount of space that is used by children of this dataset, which would be freed if all the dataset's children were destroyed. The property abbreviation is `usedchild`.
 - `usedbydataset` – Identifies the amount of space that is used by this dataset itself, which would be freed if the dataset was destroyed, after first destroying any snapshots and removing any reservation. The property abbreviation is `usedds`.

- **usedbyrefreservation** – Identifies the amount of space that is used by a refreservation set on this dataset, which would be freed if the refreservation was removed. The property abbreviation is `usedrefreserv`.
- **usedbysnapshots** – Identifies the amount of space that is consumed by snapshots of this dataset. In particular, it is the amount of space that would be freed if all of this dataset's snapshots were destroyed. Note that this is not simply the sum of the snapshots' used properties, because space can be shared by multiple snapshots. The property abbreviation is `usedsnap`.

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

```
used property = usedbychildren + usedbydataset +
usedbyrefreservation + usedbysnapshots
```

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

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

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

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

- **Using ZFS ACL Sets** – The Solaris 10 10/09 release provides the ability to apply NFSv4–style ACLs in sets, rather than apply different ACL permissions individually. The following ACL sets are provided:
 - `full_set` = all permissions

- `modify_set` = all permissions except `write_acl` and `write_owner`
- `read_set` = `read_data`, `read_attributes`, `read_xattr`, and `read_acl`
- `write_set` = `write_data`, `append_data`, `write_attributes`, and `write_xattr`

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

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

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

- “The Solaris ZFS File System” in *Solaris 10 What's New*
- “File-System Monitoring Tool” in *Solaris 10 What's New*
- “Improved Device in Use Error Checking” in *Solaris 10 What's New*

`nss_ldap` shadowAccount **Support**

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

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

Sun Validation Test Suite 7.0 Patch Set 6

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

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

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

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

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

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

Installation Enhancements

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

Turbo-Charging SVR4 Packaging

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

System Resources Enhancements

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

Zones Parallel Patching

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

For more information, see the following:

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

PVIO Drivers in Solaris 10 Guest Domains

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

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

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

Device Management Enhancements

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

iSCSI Initiator SMF Service

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

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

LSI 6180 Controller Support in Solaris MPxIO

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

System Performance Enhancements

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

Callout Subsystem Scalability

The callout subsystem is redesigned to include the following features:

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

Driver Enhancements

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

Enhanced 10 Gigabit Ethernet Performance

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

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

The ixgbe driver on x86 systems includes the following enhancements:

- TCP transmit throughput is improved to almost 100% for 8 or more connections
- TCP receive rates are 10Gb line rate for 8, 32, 100, 400, and 1000 connections
- UDP transmit maximum throughput doubles to 10Gb line rate
- Ping pong data rates are improved from 2x to 3x as the message size increases from 64 bytes to 512 bytes

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

InfiniBand Feature Improvements

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

- **InfiniBand Host Channel Adapter (HCA)** – The Solaris 10 10/09 release includes a significantly enhanced InfiniBand driver for the Mellanox ConnectX HCA. The InfiniBand driver enables InfiniBand protocols to operate over both Double Data Rate (DDR) and Quad Data Rate (QDR) InfiniBand fabrics. The driver is also integrated into the Solaris FMA framework for fault management and the driver supports relaxed ordering on SPARC systems.
- **InfiniBand Transport Framework (IBTF)** – The Solaris 10 10/09 release includes a significantly improved IBTF implementation that provides enhanced support for running RDMA-based InfiniBand protocols in Solaris. InfiniBand for SPARC now supports both the PCI Dynamic Reconfiguration (DR).
- **Internet Protocol over InfiniBand (IPoIB)** – The Solaris 10 10/09 release includes a significantly improved IPoIB driver (ibd) supporting Internet RFCs 4391 and 4392. The IPoIB driver in the Solaris 10 10/09 release supports the User Datagram (UD) mode of operation, IPv4 and IPv6 addressing, and takes advantage of hardware offloads in the ConnectX HCA for improved throughput at lower CPU utilization. IPoIB-UD enables the use of any TCP/IP application protocol such as SSH, HTTP, FTP, NFS, and iSCSI over both Double Data Rate (DDR) and Quad Data Rate (QDR) InfiniBand fabrics. The new IPoIB driver for both SPARC and x86 platforms offers a significant performance boost over the previously available driver.
- **Sockets Direct Protocol (SDP)** – The Solaris 10 10/09 release includes a significantly improved SDP driver and sockfs implementation. SDP is a transport protocol layered over the Infiniband Transport Framework (IBTF). SDP is a standard implementation based on

Annex 4 of the Infiniband Architecture Specification Vol 1. The SDP protocol provides reliable byte-stream, flow controlled two-way data transmission that is similar to the Transmission Control Protocol (TCP). InfiniBand programmers use SDP through the `Libsdp` C library that supports a sockets-based `SOCK_STREAM` interface to application programs. The SDP protocol supports graceful close, IPv4 and IPv6 addressing, the connecting/accepting connect model, out-of-band (OOB) data and common socket options. The SDP protocol also supports kernel bypass data transfers and data transfers from send-upper-layer-protocol (ULP) buffers to receive ULP buffers.

- **Reliable Datagram Sockets (RDS)** – The Solaris 10 10/09 release includes an improved RDSv1 driver certified for use with Oracle RAC (Real Application Clusters) 10gR2.
- **User-Level Direct Access Programming Library (uDAPL)** – The Solaris 10 10/09 release includes an updated uDAPL over InfiniBand API that conforms to the latest Direct Access Transport (DAT) Collaborative uDAPL 1.2 specification.

Driver for LSI MPT 2.0-Compliant SAS2.0 Controllers

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

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

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

x86: Broadcom NetXtreme II Gigabit Ethernet Driver

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

x86: Interrupt Remapping Support for Intel Vt-d

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

x86: SATA Tape Devices Supported by the AHCI Driver

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

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

Sun StorageTek 6Gb/s SAS PCIe RAID HBA Driver

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

Some of the supported RAID features include:

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

Intel 82599 10Gb PCI Express Ethernet Controller

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

Intel 82598 10Gb PCI Express Ethernet Controller

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

Freeware Enhancements

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

NTP Version 4.2.5

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

PostgreSQL

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

Samba

The Solaris 10 10/09 release supports Samba 3.0.35.

