

ORACLE SOLARIS 11 11/11 – WHAT'S NEW

#1 ENTERPRISE OS. BUILT FOR CLOUDS.

KEY FEATURES

- Automated Installer
- Image Packaging System
- Oracle Solaris Zones
- ZFS File System

KEY BENEFITS

- Simplified Administration
- Designed-in Virtualization
- Scalable Data Management
- Advanced Protection

Oracle Solaris 11 delivers groundbreaking features for fast, secure and reliable deployments in large-scale cloud environments and enterprise data centers.

Introduction

Oracle Solaris is the #1 enterprise operating system providing industry leading availability, security and performance on both SPARC and x86 systems. Oracle Solaris 11 has been tested and optimized for Oracle hardware and software, and is an integral part of Oracle's combined hardware and software portfolio.

This document covers some of the exciting new features that are available in Oracle Solaris 11 11/11 since the release of Oracle Solaris 11 Express. Make sure to also check out the new features that were introduced with Oracle Solaris 11 Express to see the complete list of features added since Oracle Solaris 10.

Download Oracle Solaris 11 today!

Installation

Automated Installer

The Automated Installer is a new modern enterprise-class installation framework for automated system provisioning. Automated Installer reduces complexity with improved integration with other Oracle Solaris technologies, helping to reduce up-front and ongoing costs of deployment. Using a network installation service, systems can be installed according to a specified installation manifest, which details system configuration, what software should be installed, and any virtualized environments that should also be provisioned.

The Automated Installer is integrated with the new package management framework, Image Packaging System (IPS) – once a system has bootstrapped a minimized operating system, installation will progress by installing software over the network from software package repositories.

Unlike Jumpstart, included in Oracle Solaris 10, Automated Installer can provision Oracle Solaris Zones as part of its base functionality. Administrators specify what zones they want to create as part of the installation manifest, and these are provisioned during the initial system reboot after the base operating system has been installed. Only Oracle Solaris 11 zones can be provisioned at this time.

Automated Installer installation images are also directly bootable, providing an easy way of installing Oracle Solaris 11 without setting up the network installation service – simply boot off the CD, specify an install manifest that exists on the network (or use the default manifest provided on the media) and the system will automatically provision itself.

Automated Installer provides an intuitive installation service management interface, allowing administrators to manage different installation services across different architectures for improved change control within the data center. New installation services can be created, deleted, and updated with minimal effort.

In enterprise environments with large scale deployments of Oracle Solaris using the Automated Installer, it is necessary to be able to derive new installation manifests from existing ones so as to allow for subtle variations in hardware attributes, software profiles, or otherwise. Derived manifest functionality allows administrators to dynamically change parameters of an Automated Install manifest through scripting using the `aimanifest(1M)` command.

Jumpstart Migration Utility

Oracle Solaris 11 includes the ability to convert Oracle Solaris 10 Jumpstart rules and profiles to Automated Installer manifests. This utility makes a best effort to translate those Jumpstart keywords that can be translated to the AI context and is not intended to create a one to one equivalence with JumpStart. Administrators wishing to use the `js2ai(1)` command line utility must install the `pkg:/install/js2ai` package.

Interactive Text Installations

The Oracle Solaris 11 interactive text installer provides users with the opportunity of installing systems without a graphical display. The installer will step through a process similar to the graphical Live Media installation, but only install a basic selection of software more suitable for server deployments. For example, it will not install components such as a graphical desktop environment, audio or wireless networking drivers, though these can be added later using package management tools if desired.

Live Media Installation

The Oracle Solaris 11 Live Media, available only for x86 based systems, provides users with an opportunity to explore a complete Oracle Solaris 11 environment without the requirement of installing it onto a system by loading the operating system into RAM. Once evaluated, customers can start a graphical installer to install the operating system. The graphical installer installs a fixed selection of software with minimal configuration, including a full desktop environment. With the inclusion of the GNU Partition Editor on the Oracle Solaris 11 Live Media, users can resize, create or delete disk partitions and file systems prior to installing the operating system.

Distribution Constructor

The distribution constructor is a command-line tool for building pre-configured, bootable customized Oracle Solaris 11 installation images for x86 and SPARC. Using a manifest description, administrators can customize target disks, software package selection, and basic system configuration to create a set of installation media that act as golden install images in the data center.

Packaging

Image Packaging System (IPS) is a new network based package management system included in Oracle Solaris 11. It provides a framework for complete software lifecycle management including installation, upgrade and removal of software packages. Integrated with the ZFS file system, IPS ensures a safe system upgrade with ZFS Boot Environments by applying system updates to a clone of the file system.

Software is installed from network based package repositories with full automatic dependency checking – any software that is required is automatically installed or updated. Should anything unexpected happen, administrators can quickly check the integrity of software packages and fix any issues, or quick boot into an older environment to minimize system downtime. IPS introduces user friendly package names, allowing administrators to quickly browse and search for packages using the command line utilities or graphical Package Manager. Additionally,

SOLARIS

IPS provides the ability to lock down individual packages on the system to ensure that software critical to business needs is not updated.

The Oracle Solaris 11 fast reboot feature means that system boot time is significantly faster by allowing certain system and firmware checks to be bypassed, helping also to minimize system downtime. The SMF service property that specifies a fast reboot is enabled by default on x86 systems. Administrators can enable a fast reboot on SPARC by modifying the property as follows:

```
# svccfg -s system/boot-config:default setprop \  
  config/fastreboot_default=true  
# svcadm refresh system/boot-config:default
```

IPS package repositories ensure simplified software delivery with a centralized architecture for managing multiple versions of software for different architectures. Administrators can control access to different software package repositories, or mirror existing repositories locally for network restricted deployment environments. A standalone 'on-disk' format for IPS packages allows administrators to install individual packages directly from an archive, rather than through package repositories in situations where network-based access is not possible or undesirable.

IPS is integrated with Oracle Solaris Zones, allowing administrators to easily update a system and any virtual environments. Each non-global zone can be installed with a different selection of software packages independently of others. Linked images, however, ensure that software versions are kept in sync between the global zone and all non-global zones and preserve the integrity of the overall system at all times. Software in the non-global zone is installed through system repositories that communicate with the global zone, caching any software content that has already been installed to ensure a faster, secure and more efficient process across all non-global zones.

SVR4 Package Compatibility

While IPS is the default package management framework in Oracle Solaris 11, compatibility for older SVR4 packages has been provided to allow administrators to install legacy software packages with `pkgadd (1M)`. However, the legacy Oracle Solaris 10 patching tools are not available on Oracle Solaris 11. Administrators wishing to patch SVR4 packages must un-install and re-install these packages.

System Configuration

With the introduction of the new installation and packaging technologies, the Service Management Facility (SMF) is an important part of the system configuration and software package installation architecture. As part of the installation process, different SMF services are activated on first reboot to apply different parts of the system configuration profile. Likewise, during a software package install, SMF services can be activated to apply configuration or refresh configuration caches as an alternative to direct post-install scripts. These changes have been made to ensure that configuration is applied in a more reliable and repeatable fashion, and ensure a more seamless transition during any system upgrade.

Layering of the SMF Repository

The SMF repository has been modified in Oracle Solaris 11 to allow better control over administrative customizations for service and system configuration, their preservation during a system update when new system configuration values are provided through manifests and improved auditing of state changes. The repository is assembled through 'layers' from a combination of current state, administrative customization through profiles, and default values as specified by manifests imported into the system. Four layers have been added in decreasing

SOLARIS

priority order – admin for changes made by interactive use of SMF commands or libraries, site-profile for any values provided in the site profile directory `/etc/svc/profile/site`, system-profile for any values provided in the system profile locations `/etc/svc/profile/generic.xml` and `/etc/svc/profile/platform.xml`, and finally manifest for any values from system manifest locations `/lib/svc/manifest` or `/var/svc/manifest`.

System Configuration in the SMF repository

Some basic system configuration has been migrated into the SMF system configuration repository with Oracle Solaris 11, as part of a wider move to reduce the number of configuration files that are located in `/etc` and to improve the management of configuration during system upgrade. System configuration can be applied during the Automated Installer process using system configuration profiles that are consumed when a series of SMF services are enabled upon first reboot.

Name service configuration has been migrated into the SMF configuration repository. Among changes to system configuration include the addition of a new service `svc:/system/name-service/switch` which manages SMF configuration that was formerly in `/etc/nsswitch.conf`, and an existing service `svc:/network/dns/client` which manages SMF configuration that was formerly in `/etc/resolv.conf`. A new utility, `nscfg(1)`, has been provided to import and export name service configuration into and out of the SMF repository and allows legacy files such as `/etc/nsswitch.conf` and `/etc/resolv.conf` to regenerate from SMF configuration for backwards compatibility.

Previously `nodename` (the system hostname) and `defaultdomain` (a host's domain name for direct use by the NIS naming service) configurations have been stored in `/etc/nodename` and `/etc/defaultdomain` respectively. These have now been migrated into the SMF configuration repository, `nodename` and `defaultname` configurations have been moved into the SMF services `svc:/system/identity:node` and `svc:/system/identity:domain`. The configuration files in `/etc` will be migrated automatically upon start up.

The default locale and time zone configuration, as stored previously in `/etc/default/init` have been migrated into the SMF configuration repository, with a new service `svc:/system/environment:init` for managing this configuration. `/etc/default/init` will be automatically regenerated from SMF configuration for backwards compatibility.

Driver configuration files (`driver.conf`) have moved into `/etc/driver/drv`. At boot time, the system checks for a configuration file in `/etc/driver/drv` for that driver. If found, the system will automatically merge the vendor-provided configuration with any local changes. `prtconf(1M)` has a new `-u` option that can be used to list both vendor and administrative driver configurations.

Resetting System Configuration

A new utility, `sysconfig(1M)`, for unconfiguring and reconfiguring an existing Oracle Solaris 11 system has been added to replace the legacy `sys-unconfig` and `sysidtool` utilities. This tool can be used within the global zone and non-global zones to configure a set of pre-defined groupings such as system identity, network, users, naming services and location/time zones. The tool can be run interactively, using the System Configuration Interactive Tool, or in a hands-off manner using a system configuration profile.

SMF and FMA Notifications

SOLARIS

New to Oracle Solaris 11 is the ability to notify administrators of SMF service state changes and FMA fault management events. Administrators can configure SNMP trap notifications and SMTP email notifications to watch certain events or services. Notifications can also be sent through to Oracle with Automated Service Requests (ASR) notifications, providing automatic telemetry to customers who have an active Oracle support agreement.

Virtualization

Oracle Solaris Zones provides built-in secure and isolated runtime virtual environments in which to deploy enterprise applications. With Oracle Solaris 11, Oracle Solaris Zones has been more tightly integrated into the operating system – zones are easier to create and manage, more flexible and functional, and provide a great level of resource management and monitoring.

Oracle Solaris 10 Zones

Oracle Solaris 10 Zones provide the ability to run Oracle Solaris 10 environments within a zone running on Oracle Solaris 11. For users already running their applications either in zones or on bare metal on Oracle Solaris 10 systems, virtual-to-virtual (v2v) and physical-to-virtual (p2v) tools have been provided to help migrate these environments to Oracle Solaris 11. Oracle Solaris 10 Zones provide a proven, tested and fully supported option for quick adoption of Oracle Solaris 11 allowing administrators to benefit immediately from all the new features available while providing an easy application migration path.

Zone Physical-to-Virtual and Virtual-to-Virtual Pre-flight Checker

The process of consolidating a system into a zone is referred to 'physical to virtual conversion' or p2v. A new pre-flight checker utility, `zonep2vchk (1M)`, has been added to Oracle Solaris 11 which can be run on the physical system prior to performing the p2v process to provide information and identify any issues in advance. This utility can also generate a proposed zone configuration based on the source system configuration and analyze any applications running on the source system for potential issues.

NFS Server in a Zone

In previous versions of Oracle Solaris, establishing an NFS share was not supported in non-global zones as it required a privilege to be granted that is prohibited by the Oracle Zones security model. New to Oracle Solaris 11, NFS servers are now supported in a non-global zone. Administrators who wish to disable shares within a zone can add `PRIV_SYS_SHARE` to the zone's set of prohibited privileges.

Exclusive-IP Zones by Default

Exclusive-IP zones provides administrators the ability to assign a separate IP stack per zone, where each zone has the flexibility to configure IP within that stack completely separate to other zones. This meant that administrators could easily observe network traffic per zone and apply individual network resources. However, in previous versions of Oracle Solaris this was dependent on the number of physical NICs (network interface controllers) an administrator had per system. The addition of network virtualization provides administrators enhanced flexibility with regard to managing zones, without the restrictions of physical network hardware. Newly created zones in Oracle Solaris 11 will be exclusive-IP zones with a VNIC (virtual network interface controller) named `net0` whose underlying lower link will be automatically selected at boot time. Shared-IP zones are also still available with Oracle Solaris 11.

Automatic VNIC Creation for Zones

For most zone deployments, administrators have some simple requirements regarding

SOLARIS

networking configuration consisting of basic IP connectivity. In order to facilitate a good user experience, a temporary VNIC is now automatically created for an exclusive-IP non-global zone. The VNIC is created when the zone boots and deleted when the zone halts, and will be created within the non-global zone's datalink namespace. This aids administrators wishing to provision a zone without having to learn the details of network configuration and topology. Administrators wishing to assign a pre-existing datalink to an exclusive-IP zone can still do so during zone configuration.

The current method of adding a datalink to an exclusive-IP zone is to add a net resource via `zonecfg (1M)`. To distinguish between the assignment of a physical resource via the existing net property and the automatic creation of a resource (a VNIC) at boot time, a new `anet` resource is being introduced as follows:

```
# zonecfg -z myzone
zonecfg:myzone> set ip-type=exclusive
zonecfg:myzone> add anet
zonecfg:myzone:anet> set lower-link=nxge0
zonecfg:myzone:anet> end
```

Administer Network Flows within Non-Global Zones

Network flows is a key element of service virtualization. Administrators can use flows to achieve bandwidth and priority control based on IP addresses, subnets, transport protocols and ports. New with Oracle Solaris 11 is the ability to administer network flows using `flowadm (1M)` and `flowstat (1M)` from within exclusive-IP non-global zones.

Delegated Administration

With Oracle Solaris 11 the administration of Oracle Solaris Zones becomes much more flexible. You have the ability to delegate common zone administration tasks for specific zones to different administrators using Role-Based Access Control (RBAC). With delegated administration, for each zone, a user or set of users may be identified with the permissions to login, manage or clone that zone. These specific authorizations will be interpreted by the appropriate commands running in the global zone to allow access at the correct authorization level to the correct user.

Zone Boot Environments

Boot Environments are also integrated with Oracle Solaris Zones, such that all non-global zone root file systems are a ZFS dataset known as Zone Boot Environment datasets (ZBE). When a new boot environment is created, by cloning an existing one, the base boot environment's zones are also cloned into the new boot environment. Support has been added to `beadm (1M)` to allow for administrators to manage nested boot environments, also known as non-global zone boot environments. Nested boot environments introduce a bootable/not-bootable distinction for a boot environment inside a non-global zone. A nested boot environment is considered non-bootable if it is not associated (via a ZFS user property) with the currently active global zone BE.

Improved Zones Dataset Layout

Oracle Solaris 11 introduces an improved ZFS dataset and boot environment layout for zones, such that non-global zones now mimic the same layout as the global zone. This gives administrators a much more consistent view across global and non-global zones. The inclusion of ZFS dataset aliases for zones help hide the portion of the ZFS dataset hierarchy that is not manageable from the non-global zone. Additionally, administrators can also now choose to store their data in ZFS datasets that aren't boot environment datasets.

Immutable Zones

Immutable Zones, or read-only root for zones, adds support for read-only file systems for non-global zones using Mandatory Write Access Control (MWAC), a policy enforced on non-global zones by the global zone. While the default for any new zone creation is to have a writeable root dataset, administrators can now use two new zone properties to define this behavior and protect the on-disk configuration of a zone.

Shutdown Zones Cleanly with `zoneadm(1M)`

`zoneadm(1M)` now supports the ability to cleanly shutdown a zone using the `zoneadm` shutdown command line option. Administrators previously had to login to the zone and perform a shutdown command, or use the `zoneadm halt` command to signal all processes and abruptly halt a zone.

Monitoring a Zone with `zonestat(1)`

The introduction of `zonestat(1)` in Oracle Solaris 11 greatly facilitates the observation of system resources consumed by Oracle Solaris Zones. More specifically, administrators can observe memory and CPU utilization, utilization of resource control limits, total utilization and per-zone utilization break-down over specified time periods. For zones that have been configured to use exclusive IP, administrators can also view a wide range of information including a high level summary with per zone breakdown, network device utilization with per zone breakdown, or detailed utilization on datalinks, virtual links and zones.

Zones Statistics Library, `libzonestat`

3rd party application developers can now use a new public C library, `libzonestat(3LIB)`, to obtain zone related resource usage statistics in a programmable way. `libzonestat` reports system wide and per-zone utilization of physical memory, virtual memory, CPU resources and networking.

Security

Role Authentication

The traditional UNIX root account is now a role by default in Oracle Solaris 11. Authorized users can assume the root role rather than directly logging into a root user account. During installation, the first user account is assigned the root role. This feature extends Oracle Solaris Role Based Access Control (RBAC) and authorized non-root users to complete tasks and scripts with superuser privileges. Added with Oracle Solaris 11 is the ability to specify whether to use the role password or user password when a user wants to assume a given role. Administrators can specify either `user` or `role` for the `roleauth` keyword. If `roleauth` is not specified, `role` is implied. Any newly created roles will be `user` by default. Additionally, `sudo(1M)` is also available to provide a familiar method for executing commands with privilege.

Trusted Platform Module

A Trusted Platform Module (TPM) chip is a hardware device usually attached directly to the motherboard of a computing platform that is intended to provide protected storage and protected capabilities on an inexpensive component with restricted resources. Oracle Solaris 11 includes driver support for TPM as specified by the Trusted Computer Group (TCG) 1.2 specification for compliant TPM devices, a TSS software stack (and a PKCS#11 provider for the Oracle Solaris Cryptographic Framework) that uses TSS to provide mechanisms for doing cryptographic operations on the secure device, and administrative tools for managing the TPM and PKCS11 provider.

Labeled IPsec

When labeled processes in a multilevel secure operating system, such as Oracle Solaris Trusted Extensions, communicate across system boundaries, their network traffic needs to be labeled and protected. Traditionally, this requirement is met by using physically separate network infrastructure to ensure that data belonging to different labeled domains stays in separate physical infrastructures. Labeled IPsec/IKE, new to Oracle Solaris 11, enables customers to reuse the same physical network infrastructure for labeled communications by transferring labeled data within separate labeled IPsec security associations, removing the need for redundant and expensive physical network infrastructure.

IPsec Supports AES GMAC Cryptographic Algorithm

Oracle Solaris 11 adds support for AES GMAC a cryptographic algorithm that implements the data integrity of AES Galois/Counter Mode (AES GCM), but without actually encrypting the data. This mode of operation is useful to administrators who don't want the encryption performance penalty, or require network data to be unencrypted for auditing purposes.

New Kerberos DTrace Providers

A new DTrace USDT provider for Kerberos has been added with Oracle Solaris 11 that provide probes for Kerberos messages (Protocol Data Unit). The probes are modeled after the Kerberos message types described in RFC4120.

Trusted Extensions Enhancements

To enable greater flexibility and security, Trusted Extensions now enables per-label and per-user credentials allowing administrators to require a unique password for each label. This password is in addition to the session login password, and thus allowing administrators to set a per-zone encryption key for each label of every user's home directory.

Oracle Solaris 11 introduces a new command, `tnctfg(1M)`, allowing administrators to create, modify and display configuration of networking properties related to Trusted Extensions, allowing administrators to label network packets received from remote hosts.

Trusted Extensions has also added support to explicitly set security labels on ZFS datasets ensuring that ZFS file systems for a specific security label cannot be mounted on a zone of a different label, and thus inadvertently upgrade or downgrade the classification of data.

Support for ssh X.509 Certificate Extension

Oracle Solaris 11 introduces support for ssh X.509 certificate extension that allows for the use of a set of trusted certificates instead of public key distribution during public key authentication. If configured, administrators will no longer need to respond to client-side questions of host authenticity nor need to populate `~/.ssh/authorized_keys` with user public keys on the server-side.

Oracle Solaris Cryptographic Framework

To meet more stringent government standards, the Oracle Solaris Cryptographic Framework now supports the NSA Suite B algorithms. Additionally, the Oracle SPARC T4 processor supports the AES CFB mode used by table space encryption features of the Oracle Database Advanced Security Option. This is closely related to the ability of the Oracle Solaris Cryptographic Framework to provide acceleration through on-board cryptographic mechanisms on both SPARC and Intel chips. The Oracle Solaris Cryptographic Framework also includes support for the Intel Advanced Encryption Standards – New Instructions (AES-NI).

To help with the complex task of managing keys, the Oracle Key Management System can now be used for AES key storage using the new `pkcs11_kms` plugin for the Oracle Solaris

Cryptographic Framework. This mechanism can be used by any PKCS#11-aware application.

In-Kernel `pfexec`, Forced and Basic Privileges

An in-kernel `pfexec` implementation, new to Oracle Solaris 11, is used to execute administrative commands requiring a higher privilege level. A new process flag is used to specify that all subsequent program executions are subject to RBAC policy. The flag is set at the first invocation of any of the complete set of profile shells (`pfsh(1)`, `pfcs(1)`, `pfksh(1)`, `pfksh93(1)`, `pfbash(1)`, `pfrcsh(1)`, `pfzsh(1)`, `pfexec(1)`) and inherited by child processes. This feature eliminates the need to modify shell scripts to invoke `pfexec` or profile shells. Another application of this feature is to limit the set of privileges given to programs with `setuid` to root. Processes that require the `setuid` mechanism traditionally ran with all privileges. Now they execute with only those privileges that are specified in their entry in the forced privileges rights profile, significantly reducing their potential to be an attack vector against the system. Additionally, Oracle Solaris 11 adds three new “basic” privileges (`file_read`, `file_write`, and `net_access`) beyond the five that exist in Oracle Solaris 10. These new privileges satisfy a long-standing customer need to restrict read, write and outbound network access.

Networking

The networking stack in Oracle Solaris 11 has undergone significant re-architecture to unify, simplify and enhance observation and interoperability of network interfaces and features. The introduction of a new network driver framework, GLDv3, provides VLAN, link aggregation and the ability to support MAC layers other than Ethernet (IP tunnels, Wi-Fi, Infiniband) providing flexible network administration with `dladm(1M)`. Enhancements to `dladm` also include the ability to allow links to be renamed, including non-GLDv3 links, and set properties for NIC drivers using a common command.

Network Virtualization and Resource Management

Oracle Solaris 11 introduces built-in network virtualization and resource management providing more effective sharing of network resources and enhancing the ability to consolidate server workloads. Using the basic building blocks of Virtual Network Interface Controllers (VNICs), virtual switches and interconnects, Virtual LANs (VLANs), and routing and firewall functionality, it is possible to consolidate an entire distributed computing environment on a single system for prototyping, testing and deployment scenarios without the restriction of physical network devices attached to the system.

Network resource management allows organizations to meet quality of service goals for networking. These management capabilities allow setting bandwidth limits on NICs/VNICs, and assigning CPU resource limits for servicing the NICs/VNICs. This enables organizations to create OS-enforced network sharing policies.

The new architecture has many features to allow it to work more effectively with the latest generation of intelligent NICs while still maintaining compatibility with older NICs. Highlights of the new architecture include the ability to shift from interrupt-driven to polling in high volume traffic meaning more efficient handling of network traffic, the ability to offer Quality of Service capabilities without adding any additional overhead and the ability to mitigate the effects of a denial of service attack by dealing with packets at the NIC level.

The power and flexibility of built-in network virtualization has been tightly integrated with Oracle Solaris Zones, allowing each non-global zone to have its own exclusive IP stack using VNICs without the limitation of dedicating a physical network interface (NIC) to a zone. In addition, this capacity is also extended to Oracle Solaris 10 Zones.

Manual and Automatic Networking

Support for network profiles has been added with Oracle Solaris 11 to give a much more seamless and persistent network experience. Managed by a single SMF network service, `svc:/network/physical:default`, administrators can switch between automatic and manual networking by enabling `Automatic` or `DefaultFixed` network configuration profile, or by creating their own profile, through the `netadm(1M)` and `netcfg(1M)` command line utilities. Automatic networking, the default for the Live Media install, discovers and connects to networks (both wired and wireless) depending on network conditions and is especially useful for laptops when network mobility is desired.

Default Names for Datalinks

The driver based naming scheme for Oracle Solaris has been changed to adopt a generic, 'vanity', name by default so as to separate high-level configuration out from low-level hardware specifics. This has the benefit of making hardware replacement and configuration migration much easier, such as providing network configuration in an Automated Installation system configuration profile. As an example, a datalink with a traditional `e1000g0` naming is now replaced by a generic `net0` name, based on the relative physical location of network devices in the system. Administrators can choose to revert to the traditional scheme if necessary.

Changing MAC Address with `dladm(1M)`

Support for changing the MAC address has been added to `dladm(1M)`. Unlike changing the MAC address using `ifconfig(1M)`, this is a persistent change across reboots, and changes the primary MAC address used by all current and future MAC clients of the underlying data link.

InfiniBand, Enabled and Optimized

The InfiniBand stack has undergone some significant improvements in Oracle Solaris 11, including improved support for Sockets Direct Protocol (SDP) allowing the transparent redirection of TCP/IP usage to SDP and the efficiencies it brings, and the addition of the RDSv3 protocol providing better performance and observation for Oracle RAC databases.

Oracle Solaris tools and utilities have been also updated to work with Sockets Direct Protocol (SDP) and take advantage of high-performance RDMA network features such as zero-copy data transfers. The list of utilities updated includes `netstat`, `truss`, `files`, `mdb` and `kldb`. Additionally, support for SDP within a non-global zone (both shared and exclusive IP stacks) is now provided with Oracle Solaris 11.

Registration of Virtual LANs

New to Oracle Solaris 11 is the ability for broadcasting VLAN ID information to the network fabric. Network virtualization enables creating software Virtual Network Interfaces (VNICs) that are associated with real, physical networking ports and associating VLAN ID for outgoing and incoming traffic. In order for this traffic to be forward over an actual fabric, for example a network switch, the fabric must be configured to accept traffic from the specific VLANs of each individual VNIC. This is especially useful in a cloud environment of many virtual machines, where administrators now have the ability to create multiple VNICs and VLANs and have the network fabric autoconfigure itself.

Link Layer Discovery Protocol Support

Support for Link Layer Discovery Protocol (LLDP) has been added to Oracle Solaris 11. LLDP is a one-way link layer protocol that allows an IEEE 802 LAN station to advertise the capabilities and current status of the system to other stations attached to the same LAN. The

SOLARIS

`lldpadm(1M)` utility has been added to enable or disable an LLDP agent on a physical datalink.

New Sockets Architecture

The socket implementation has been re-written for Oracle Solaris 11 and no longer uses STREAMS. Historically Oracle Solaris has supported streams-based sockets, and with the move to the new architecture there are significant performance improvements along with a simplified developer interface for adding new socket types.

Load Balancing

Oracle Solaris 11 includes an integrated L3/L4 load balancer. This capability can be complimentary to existing higher layer load balancing solutions from different ISVs. The addition includes stateless DSR and NAT operation modes on a variety of load balancing algorithms, a command-line and configuration API to configure various features as well as view statistics and other configuration details.

Link Protection

In many virtualized setups today, it is common for the host administrator to grant exclusive access of a physical link or a Virtual NIC to a guest VM. This enables guests to benefit from traffic isolation and improved performance. The downside is that guests are allowed to generate any type of packet, even harmful ones, to the network. Link protection is a new mechanism for preventing potentially malicious or misbehaving guest VMs from sending harmful packets to the network. This feature provides protection against these basic threats: IP, DHCP, MAC and L2 frame spoofing. Unlike a traditional firewall, link protection does not support inbound filtering or customizable filtering rules. For users with such requirements, a firewall should be used instead such as Oracle Solaris IP Filter, `ipf(1M)`.

Bridging and Tunneling

Bridging is a general layer two (L2 or datalink) technology that is used to connect together separate L2 subnetworks, allowing communication between attached nodes as if only a single subnetwork were in use. Basic ethernet bridging support has been added to Oracle Solaris 11 using the Spanning Tree Protocol (STP, IEEE 802.1D-1998) and TRILL protocol. IP tunneling functionality has been reimplemented in Oracle Solaris 11, delivering a generic LAN driver (`iptun`) that implements IP tunnel links atop which IP interfaces can be plumbed and managed through `dladm(1M)`. With this new architecture, tunnel links gain functionality common to other links, including link vanity naming, link-layer observation using transitional tools such as `wireshark(1)` and `snoop(1M)`, and the assignment of tunnel links to exclusive stack non-global zones.

IP Observability

Oracle Solaris 11 improves the area of IP observability, allowing the developer or administrator to use common packet sniffing tools such as `wireshark(1)` and `snoop(1M)` to view all IP traffic sent on real and virtual paths. Now all traffic can be observed at the IP layer, including to and from an Oracle Solaris Zone. In addition, Oracle Solaris 11 includes `dlistat(1M)`, a tool to provide runtime statistics for data links allowing administrators to get a better understanding of how well their networks are performing.

IP Multipathing

IP Multipathing (IPMP) provides transparent redundancy for IP level communications between the application running on a system, and the first router in the communication path to the outside world. IPMP allows creating multiple paths to that first router so that port, NIC, cable or switch failures will not impact any connections. For high availability applications

SOLARIS

IPMP will assign one of the interfaces to the IP address and continually monitor the underlying interfaces to ensure a connection is maintained. If IPMP detects that the IP interface being used has failed it will use an alternate working IP interface. Applications do not need to be aware they are running on a system managed by IPMP. For Oracle Solaris 11 Express 2010.11, IPMP has undergone a significant re-architecture to improve network administration and network observation.

IPMP Administration through `ipadm(1M)`

Continuing the consolidation of network management tools, IPMP can now be administered using the `ipadm(1M)` command line utility with a number of new sub-commands to support the creation of IPMP interfaces and groups.

IPMP Transitive Probing for Less Restrictive Availability

The addition of a new failure detection mode for IPMP, Transitive Probes, provides an alternate probe-based failure detection between the system and first-hop router without the restriction of requiring additional test IP addresses as is the case for ICMP probe detection. This has the benefit that administrators can now easily deploy IPMP groups within a virtualized environment, particularly Oracle Solaris Zones using exclusive-IP, without such a restriction. IPMP transitive probing is turned off by default, and can be enabled as follows:

```
# svccfg -s svc:/network/ipmp setprop config/transitive-probing=true
# svcadm refresh svc:/network/ipmp:default
```

Improved Availability for `in.mpathd`

`in.mpathd`, the daemon that performs failure and repair detection for IP interfaces that have been placed into an IPMP group, is now managed by an SMF service `svc:/network/ipmp` for significantly improved availability. This service, like all other services on Oracle Solaris, can now be restarted automatically should a failure occur for any reason.

I/O Enhancements to `netcat`

`netcat`, a popular utility used for network observation and debugging, has been enhanced to include a number of command line options that allow administrators to configure a number of previously hard coded values related to I/O and integration with zones, along with a number of usability additions to improve familiarity with `netcat` on other operating systems.

New FTP Server Implementation

In prior Oracle Solaris releases, the FTP server implementation was based on `WU-ftpd`. In Oracle Solaris 11 this has been replaced with `proftpd` with an improved feature set and enhanced level of security. This FTP server is already used in Oracle's ZFS Storage Appliance.

DTrace Networking Providers

The `tcp`, `udp` and `ip` DTrace providers allow administrators to trace the TCP, UDP and IPv4/IPv6 network protocols.

Storage

ZFS is the root file system on Oracle Solaris 11 offering a superior experience in terms of manageability, scalability and data integrity. ZFS presents a pooled storage model that completely eliminates the concept of volumes and the associated problems of partitions, provisioning, wasted bandwidth and stranded storage. Thousands of file systems can draw from a common storage pool, each one consuming only as much space as it actually needs. All

SOLARIS

operations are copy-on-write transactions ensuring that the on-disk state is always valid. Additionally, blocks are check-summed to prevent silent data corruption, allowing data to self-heal itself in replicated (mirrored or RAID) configurations. If one copy is damaged, ZFS detects it and uses another copy to repair it. ZFS is also at the heart of Oracle Solaris 11 software installation and management with the IPS packaging system, greatly reducing planned and unplanned downtime with safe system upgrade capability. UFS is no longer supported as a root file system, though UFS file systems can still be mounted.

ZFS Dataset Encryption

Encrypted dataset support has been added to ZFS to protect against theft of physical storage, man-in-the-middle attacks on the SAN, and to provide dataset level secured deletion. Data is encrypted at the dataset level, allowing a mix of encrypted and unencrypted datasets in the same ZFS storage pool. A single dataset has a consistent policy only allowing for encryption to be set at dataset creation time. All data and file system metadata is encrypted with a comprehensive encryption key management facility to cover different key management strategies. Encrypted root pool is not supported at this time.

Save Space with ZFS

Deduplication is a feature of modern storage platforms by which varying mechanisms are employed to reduce the amount of total data stored by eliminating and sharing common components. ZFS deduplication support has been added to Oracle Solaris 11. ZFS deduplication uses checksum based comparison of blocks with optional verification (for example with non-cryptographically secure checksums). Deduplication is performed across the entire ZFS storage pool; administrators can select if individual datasets have deduplication enabled or not. This is useful in mixed-mode environments in which some datasets have highly duplicated data (e.g. virtualized images, home directories, or email folders) and others are unique (e.g. databases). Deduplication can be used in combination with ZFS compression. However, when used in combination with ZFS encryption, deduplication of data can only occur within a single dataset or clone of that dataset as encrypted datasets by default have different data encryption keys.

ZFS Shadow Migration

Shadow migration is a tool for migrating data from an existing file system to a new file system. A "shadow" file system is created that pulls data from the original source as necessary and falls through to the native file system for reads and writes once a file has been migrated. Oracle Solaris 11 introduces a new "shadow" ZFS dataset property that can be used either for migrating a local to the machine file system or for migrating an NFS file system.

ZFS Backup with NDMP

Support for backing up and restoring ZFS volumes has been added to the Oracle Solaris Network Data Management Protocol (NDMP) service. While the NDMP service has been able to back up ZFS file systems using tar previously, the addition of ZFS support using ZFS send and receive provides the ability to back up ZFS volumes or raw partitions seamlessly, including the capability of backing up ZFS descendants of the dataset, including snapshots and clones. This new capability also provides significant performance improvements in NDMP backup and restore environments.

Temporary ZFS Mounts

A file system can now be mounted temporarily at a location other than the file system's persistent mount point by specifying the `-o mountpoint=value` option to the `zfs mount` command. This is only permitted for file systems with non-legacy mount points. This

SOLARIS

is especially useful when maintaining boot environments whose persistent mount point is `/` but cannot actually be mounted at that location for maintenance purposes because the root is already occupied.

ZFS Snapshot Alias

A convenient shorter command line option has been added with Oracle Solaris 11 to allow administrators to take a ZFS snapshot using `zfs snap`. Administrators can still use `zfs snapshot` as well.

Recursive ZFS Send

New functionality has been added to Oracle Solaris 11 to support a recursive ZFS stream using `zfs send`. A recursive stream package consists of the specified dataset and its descendents. Similar to a replication stream, a recursive stream does not include any unneeded intermediate snapshots. Additionally, administrators can now also create entirely self-contained recursive streams.

ZFS Diffs

Support for listing the differences between ZFS snapshots has been added with Oracle Solaris 11. Users with the appropriate privilege can now view what file and directory level changes have occurred between snapshots, such as files or directories added, removed, modified or renamed in a later snapshot.

NFSv4 Client and Server Migration Support

The NFS Version 4 protocol defines the way a client and a server are to handle file system migration. NFS support for Oracle Solaris 11 has been extended to ensure that a client will be extended to react appropriately when a file system has been migrated, which is generally transparent to applications. An NFS server will have the ability to quiesce the file system against changes, save in-memory state on the source, restore the state on the destination, and inform the client of the move.

Built-in CIFS for Microsoft Interoperability

Oracle Solaris 11 includes fully integrated CIFS. The Common Internet File System (CIFS) also known as SMB, is the standard for Microsoft file sharing services. The Oracle Solaris CIFS service provides file sharing and MS-RPC administration services required for Windows like behavior for interoperability with CIFS clients, including many new features such as host-based access control which allows a CIFS server to restrict access to specific clients by IP address, ACLs (access control lists) on shares, and synchronization of client-side offline file caching during reconnection. Microsoft ACLs are also supported in ZFS.

DTrace Storage Providers

A new DTrace provider, `smb`, has been added to allow administrators and developers to observe a wide range of SMB operations, before and after requests are executed. The `iscsi` DTrace provider allows administrators to trace iSCSI target activity from the server's perspective.

COMSTAR SCSI Target Framework

COMSTAR (Common Multiprotocol SCSI Target) is the software framework that enables the ability to turn any Oracle Solaris host into into a target device that can be accessed over a storage network. The COMSTAR framework makes it possible for all SCSI device types (tape, disk, and the like) to connect to a transport (such as Fibre Channel) with concurrent access to all logical unit numbers (LUN) and a single point of management. Support for a number of protocols has been added; iSCSI Extensions for RDMA (iSER) and SCSI RDMA

Protocol (SRP) for hosts that include an InfiniBand Host Channel Adapter, iSCSI, and Fibre Channel over Ethernet (FCoE). Oracle Solaris DTrace probes have also been added to COMSTAR in the SCSI Target Mode Framework (STMF) and SCSI Block Device (SBD).

Kernel/Platform Support

SPARC T4 Support

Oracle Solaris 11 adds support for the new next generation SPARC T4 processor and Oracle SPARC T-Series servers, and takes advantage of some of the unique hardware features out-of-the-box including support for ISA cryptographic hardware optimization, 2GB page size support, CPU and DRAM performance counter support, and L3 cache support. In particular, Oracle Solaris 11 provides performance gains of between 20% and 40% for various cipher and hash instructions, significant performance gains for SSL transport, and direct cryptographic acceleration support for Oracle DB (11.2.0.3) when used in conjunction with Oracle Solaris 11.

Critical Threads

Oracle Solaris provides mechanisms such as processor sets that allow an administrator to provision specific threads with the amount of resources required for optimal performance. However, these existing mechanisms require a considerable amount of time to administer and tune. Current and upcoming processor design allows the dynamic allocation of hardware resources to provide boosts in performance. Critical threads, new in Oracle Solaris 11, helps to take advantage of these new processor designs by matching a thread's hardware requirements with the amount of exclusive resources needed to induce the activation of one of these features and guaranteeing exclusive access to the specific hardware resources.

Single-root I/O Virtualization

As enterprises push for better returns through higher consolidation ratios and the virtualization of all their applications, software emulated I/O is rapidly becoming a limited factor for virtualization. The demand to virtualize I/O intensive applications such as database and technical/compute-intensive applications and move to a fully virtualized, dynamic datacenter requires an I/O architecture that can deliver near native performance, increased throughput, and flexibility. New to Oracle Solaris 11 is support for Single-root I/O virtualization (SR-IOV) framework, defining extensions to the PCI Express (PCIe) specification to allow efficient sharing of PCIe devices among virtual machines, both in hardware and software. Support for a number of SR-IOV capable platforms has also been added.

NUMA I/O

Many modern systems are based on a NUMA (Non-Uniform Memory Access) architecture, where each CPU or set of CPUs is associated with its own physical memory and/or devices. For best I/O performance on these systems, the processing associated with a device should be performed close to that device, and the memory used by that device for DMA and PIO should be allocated close to that device as well. Oracle Solaris 11 adds support for a NUMA I/O architecture, which allows for operating system resources (kernel threads, interrupts, and memory) to be placed on physical resources according to the physical topology of the machine, specific high-level affinity requirements of I/O frameworks, actual load on the machine, as well as resource control and power management policies

Intel Advanced Vector Extensions

Support for Intel's Advanced Vector Extensions (AVX) has been added to Oracle Solaris 11. AVX introduces new instructions that accelerate compute intensive vector floating point operations common in image, video and audio processing, engineering oriented applications

such as 3D modeling and analysis, scientific simulation and financial analytics. AVX will allow applications to be optimized for the latest generation Intel micro-architecture, code-named Sandy Bridge.

Dynamic Intimate Shared Memory Performance Improvements

Significant integration work has been done in Oracle Solaris 11 to improve the performance of the Oracle Database stack for Oracle Solaris systems with large memory. Speed improvements to Intimate Shared Memory (ISM) and Dynamic Intimate Shared Memory (DISM) creation, locking and destruction have resulted in up to 8x start up performance improvement for the Oracle Database. Oracle Database uses DISM in its dynamic System Global Area (SGA) capability, which forms the part of the RAM shared by all processes belonging to a single Oracle Database instance.

Suspend and Resume to RAM

Support for a limited number of platforms has been added in Oracle Solaris 11 Express to allow suspend and resume to RAM. Oracle Solaris continues to advance in the area of power management as energy efficiency in the datacenter becomes more important to drive costs down and increase utilization.

Improved Hardware Support

Oracle Solaris 11 introduces support for a number of new hardware platforms and hardware components, including support in many components for Oracle Solaris Fault Management Architecture (FMA) which provides fault tolerance and fine-grained isolation of hardware components ensuring continuous service. Additions to the FMA framework include a generic topology enumeration for platform independence, a generic hotplug framework for any hotpluggable bus and migration functionality in a virtualization environment through virtual hotplug, and support for Intel's latest generation microprocessor architecture, code named Sandy Bridge. Check out the Oracle Solaris 11 Hardware Compatibility List for a full list of hardware components that is supported with this release.

Measuring System Latency

Oracle Solaris 11 includes a port of Intel's LatencyTOP, a tool that detects latency in the system and what the cause is. Through innovative use of Oracle Solaris DTrace, latencies on the system can now be measured and fixed.

DTrace cpc Provider

The `cpc` provider gives you the ability to profile your system by many different types of processor related events. The list of events is processor specific and includes events such as cycles executed, instructions executed, cache misses, TLB misses and many more. From a high level view, the `cpc` provider is essentially the same as the `profile` provider except that, where the `profile` provider allows you to profile a system using a fixed, time based source, the `cpc` provider lets you profile by events related to processor activity.

User Environment

Popular Open Source Software Packages

Over 850 popular open source software packages have been included in the IPS package repository, including Java SE 6 and 7, GCC 4.5.2, Python 2.7, Perl 5.12, Ruby 1.8.7, PHP 5.2.17 and a complete web stack. Many of the runtime languages have been integrated with DTrace to provide unprecedented levels of observability.

Enhanced Desktop Environment

SOLARIS

Oracle Solaris 11 includes GNOME 2.30.2, an easy to use desktop environment, Firefox 6, the popular open source web browser, and Thunderbird 6, the popular email, addressbook and calendaring application. Additional desktop software can be browsed and installed using the Package Manager. For desktop and laptop users, switch on Visual Effects in the Appearances preferences dialog to turn on Compiz for some stunning visual effects that takes advantage of the latest graphics hardware. A number of development libraries have also been added to the X windows stack, including XCB and FreeGLUT.

Default \$PATH and User Familiarity

A number of existing Oracle Solaris utilities have been enhanced to include popular GNU-like command line options for user familiarity. Where possible, GNU utilities have been included in `/usr/bin` if there is no existing utilities that have a namespace clash, and otherwise been included in `/usr/gnu/bin`. Oracle Solaris 11 has also introduced new levels of familiarity with Linux and BSD operating systems with a number of common command-line arguments to existing utilities, like for example the addition of `-iname` option for `find(1)` or compression support for `tar(1)`.

Default Shell

`bash(1)` is now the default shell for new users added to the system using `useradd(1M)` command line utility if another alternative shell is not provided, and the default shell for new installations. `ksh93(1)` is used as the default system shell.

Removable Media

The user experience for removable media has been improved with Oracle Solaris 11, replacing the legacy volume management daemon framework in Oracle Solaris 10. The new framework uses a combination of Hardware Abstraction Layer (HAL) and D-Bus message passing system used in popular Linux distributions to include various aspects such as hotplug, device discovery, content recognition as well as improving usability, scalability and performance across all layers of the software stack from the device drivers to the desktop application environment.

New Sound System

A new audio subsystem for supporting audio devices has been included in Oracle Solaris 11. Supporting both current and future generation multimedia audio applications and devices, a new device driver interface and compatible Open Sound System API ensures that write new drivers or porting applications from other platforms only require a modest amount of effort.

Search for Content in Manual Pages

Oracle Solaris 11 has introduced a new capability to search for man page content with query strings with the `man -K searchstring` command, including a SMF service to automatically index man pages for search. Administrators can re-index their systems by refreshing the `svc:/application/man-index` SMF service.

Virtual Console Terminals

Oracle Solaris 11 supports switching between X session and virtual console terminals, using the `svc:/system/vtdaemon:default` and `svc:/system/console-login:vt*` services in SMF. When enabled, users can switch between sessions using hotkeys `Alt-Ctrl-F#`.

Time Slider Snapshot Management

Time Slider provides the ability for users to quickly take ZFS snapshots of their home directories automatically and manually as needed. Snapshots can be viewed graphically over

SOLARIS

time with Oracle Solaris ZFS snapshot file manager integration to identify accidentally modified or deleted files.

CUPS Printing

The Common UNIX Printing System (CUPS) has been selected as the print service on Oracle Solaris 11, replacing the LP print service. CUPS support includes a web and graphical interface to manage your printing environment. A system that is running CUPS becomes a host that can accept print requests from client systems, process those requests, and then send them to the appropriate printer. Legacy LP commands remain for compatibility reasons but wrap CUPS functionality.

libc Familiarity

A number of enhancements have been made to the Oracle Solaris C library to improve familiarity with Linux and BSD operating systems and help developers port their applications across to the Oracle Solaris platform. Among these changes is the addition of common string functions (`asprintf()`, `vsprintf()`, `getline()`, `strdupa()`, `strndup()`), compatibility for date/time format conversion functions (`ascftime(3C)`, `cftime(3C)`, `wcsftime(3C)` and `fnmatch(3C)`) and file or pathname matching (`fnmatch(3C)`).

paths.h Path Name Definitions

The introduction of `/usr/include/paths.h` for definitions of path names for Oracle Solaris and `/usr/include/sys/paths.h` for definitions of path names used by the Oracle Solaris kernel provide interfaces developers can use for well known locations in which to port their applications to the platform.

Locales and Languages

Oracle Solaris 11 supports over 200 locales. Languages supported include: Afrikaans, Albanian, Arabic, Armenian, Assamese, Azerbaijani, Bengali, Belarusian, Bosnian, Bulgarian, Catalan, Chinese – Simplified, Chinese – Traditional, Croatian, Czech, Danish, Dutch, English, Estonian, Finnish, French, German, Greek, Georgian, Gujarati, Hebrew, Hindi, Hungarian, Icelandic, Indonesian, Italian, Japanese, Kannada, Kashmiri, Kazakh, Kirghiz, Korean, Kurdish, Lithuanian, Latvian, Macedonian, Malay, Malayalam, Maltese, Marathi, Norwegian Bokmal, Norwegian Nynorsk, Oriya (Indian), Punjabi (Indian), Polish, Portuguese, Portuguese – Brazil, Romanian, Russian, Sanskrit, Serbian, Slovak, Slovenian, Spanish, Swedish, Tamil, Telugu, Thai, Turkish, Ukrainian, and Vietnamese.

Locales which support translated messages are: Japanese, Chinese – Simplified, Chinese – Traditional, Korean, French, German, Italian, Spanish and Portuguese – Brazil.

TrueType Fonts

Oracle Solaris 11 improves the choice of TrueType fonts that are available on the system. There are updated and newly added font families to support the many locales and languages covered by this release.

Contact Us

For more information about Oracle Solaris 11, visit oracle.com or call +1.800.ORACLE1 to speak to an Oracle representative.



Oracle is committed to developing practices and products that help protect the environment

Copyright © 2011, Oracle and/or its affiliates. All rights reserved.

This document is provided for information purposes only and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark licensed through X/Open Company, Ltd. 0611

Hardware and Software, Engineered to Work Together