ORACLE

SOLARIS

An Oracle Technical White Paper
January 2013

# IBM AIX to Oracle Solaris Technology Mapping Guide

Preparing for the Move to Oracle Solaris 11

ORACLE

# Chapter 1. Introduction

Many IT organizations have opted to keep their enterprise data center servers running the IBM AIX 5.3 operating system, even though IBM ended standard support for the platform in April, 2012. In many of these organizations, IT managers are hesitant to upgrade to the latest version (AIX 7.1)—enough that IBM created a special program[1] for organizations that want to remain on AIX 5.3 and run it on bare metal hardware. Unfortunately, staying on an older platform means IT organizations are unable to reap the benefits of technological innovation. Another approach is to move to an alternative platform that can deliver the scalable functionality, performance, reliability, availability, and security needed to support business priorities. Oracle's portfolio of SPARC and x86 servers running the Oracle Solaris operating system provide an obvious alternative and safe platform for running business-critical applications.

When evaluating upgrade options, IT organizations must look at AIX 7.1 and Oracle Solaris and determine which environment presents the best opportunity for optimizing infrastructure today without limiting how systems can be used to address priorities in the future. Since both AIX and Oracle Solaris are based on UNIX System V, the transition to Oracle servers running Oracle Solaris 11 is not difficult. Yet migrating to a new platform can take time and effort, particularly if the products and tools used on the new system are unfamiliar.

Aimed at technical IT managers, IT architects, and system administrators tasked with moving—or evaluating the move—to Oracle Solaris, this guide compares the key tools and technologies commonly used in AIX 7.1 environments with those used in Oracle Solaris. Different concepts, processes, and technologies that are essential to successful Oracle Solaris platform deployment are identified. For each topic area, AIX features and tools are mapped to their Oracle Solaris 11 counterparts, with discussion centering on similarities and differences in functionality to help technical staff quickly identify where appropriate, equivalent resources (product and technology information, manuals, and training) are needed to support deployment.

- Chapter 2, "Distributing, Installing, and Managing Software," identifies and compares the key tools used for software management.
- Chapter 3, "Managing Data," discusses file system availability, volume management, data backup, swap space considerations and more, as well as whether and how these technologies differ in the two environments.
- Chapter 4, "Virtualizing Infrastructure," maps the key virtualization technologies used in AIX to similar Oracle Solaris virtualization mechanisms, and highlights similarities and differences that are important to understand before deployment.
- Chapter 5, "Keeping Systems, Applications, and Services Available," compares the commonly used AIX and Oracle Solaris tools for maximizing availability.

[1] See http://www.theregister.co.uk/2012/02/17/ibm_aix_5_3_i5_v5r4_support/and http://www-01.ibm.com/common/ssi/rep_ca/6/897/ENUS612-006/ENUS612-006.PDF.

- Chapter 6, "Securing Infrastructure," relates the security mechanisms in AIX to the extensive defense-in-depth approach in Oracle Solaris.
- Chapter 7, "For More Information," provides a comprehensive list of references to more detailed information. For readers interested in the Oracle Solaris 11 product documentation, see http://www.oracle.com/technetwork/server-storage/solaris11/documentation/index.html.
- Appendix A, "Glossary," defines terms used throughout this document.

# Chapter 2. Distributing, Installing, and Managing Software

While AIX and Oracle Solaris have similar concepts for software management, the tools used are very different. This chapter discusses several aspects of software management, including software installation, packaging, updates, and upgrades. It identifies which tools are available in Oracle Solaris to perform various tasks, and explains how those tools differ from those commonly used in AIX deployments.

Table 2-1 maps software management tools in AIX 7.1 to counterpart tools in Oracle Solaris 11.

**TABLE 2-1. SOFTWARE MANAGEMENT MAPPINGS**

| TASK OR CAPABILITY | AIX | ORACLE SOLARIS |
|---|---|---|
| Software packaging model and tools | Filesets, software bundles, RPM, ISMP GUI, TUI, and command line interfaces provided | Image Packaging System (IPS); GUI and command line interfaces provided |
| Single system installations | From DVD, virtual media, Base Operating System (BOS) image or using Network Installation Management (NIM) | From DVD using Live Media (x86) or interactive text installers (SPARC, x86); USB images available for Live Media and text installer; Automated Installer (AI) and post-installation customization can be performed using IPS. |
| Automated installation of multiple systems | NIM | Automated Installer and IPS software repositories |
| Adding software packages to an existing installed system | System Management Interface Tool (SMIT) | IPS tools: `pkg install` (command line) or `packagemanager` (GUI) |
| Analyzing and applying patches | Service Update Management Assistant (SUMA) | No patching required (package updates applied instead) |
| Updating software (single system) | SUMA | IPS tools: `pkg update` (command line); `packagemanager` or `pm-updatemanager` (GUI) |
| Minimizing downtime and enabling recovery for updates | Backup and restore `alt_disk_install` | Boot Environments (BEs); `beadm` utility for BE management |
| Creating customized installation images | `mksysb` | Distribution Constructor and sample manifests |

## Clean Installations

When installing a new instance of AIX on a single system—a new and complete Base Operating System (BOS) install—system administrators typically use distribution media or perform virtual media-based networked installations. Oracle Solaris 11 offers similar installation options. Interactive installations from media using Oracle Solaris 11 Live Media for x86 DVD provide a full desktop environment, while the interactive text-based user interface creates server installations for both x86 and SPARC installations.

In addition, Oracle Solaris 11 supports a hands-off automated installation process, called Automated Installer, that is analogous to the AIX Network Installation Manager (NIM). Automated Installer relies on software repositories that loosely resemble AIX software bundles.

For more information on installation options, see the *Installing Oracle Solaris 11 Systems* manual at http://docs.oracle.com/cd/E23824_01/html/E21798/index.html

## Software Packaging Model

Developers who create custom applications should become acquainted with the software-packaging model, called the Image Packaging System[2] (IPS), in Oracle Solaris 11. Use of these facilities ensures system administrators have all the resources needed for successful application deployment.

IPS provides broad software management functionality that is similar to RPM repositories or NIM-based collections of *filesets*. Designed to support management tasks for operating system as well as application software, IPS is a comprehensive framework that spans the full software lifecycle, addressing functions such as installation, patching, upgrades, and software removal. During software package installations, IPS performs automatic dependency checking, adding any additional packages (such as libraries) that might also be required. A snapshot of the system is taken before each package installation, ensuring the system is always in a valid state and enabling a rollback to be performed in the event that package installation fails.

An IPS software package identifies all necessary installable objects in a well-defined format, specifying directories, files, links, drivers, dependencies, groups, users, and license information. IPS packages include attributes such as the package name and a brief description. A Fault Management Resource Identifier (FMRI) uniquely represents each package and consists of a publisher, package name, and version number with the scheme "pkg" as in *scheme://publisher/package_name@version:dateTtimeZ*. Since the FMRI incorporates an explicit version number and timestamp, IPS easily can determine whether a more up-to-date package release exists. Specifying the package publisher in the FMRI identifies the package developer, supplying a mechanism by which IPS can classify packages, confirm authenticity, and restrict installation.

[2] IPS is new in Orale Solaris 11. IT staff that last investigated prior versions Oracle Solaris should take a close look at IPS. While the term *package* has been retained from earlier Oracle Solaris releases, IPS is very different than the original System V packaging system and tools.

For more information on the IPS model and tools, se the *Adding and Updating Oracle Solaris 11 Software Packages manual at* http://docs.oracle.com/cd/E23824_01/html/E21802/index.html

## Managing Software on an Installed System

Similar to NIM and the repositories it manages, IPS relies on software depots (called software *repositories* in Oracle Solaris 11) to access software packages for installation and update. IPS supports DVD, CD, and file-based local repositories, as well as network-based remote repositories. Administrators can easily set up and manage local repositories to deploy packages within network-restricted and firewalled environments. The default repository for Oracle Solaris 11 is http://pkg.oracle.com/solaris/release, which is publicly available. Customers with support contracts can access the support repository at My Oracle Support (http://support.orcalce.com) to obtain packages with the latest bug fixes and updates.

For more information about setting up local repositories, see the *Copying and Creating Oracle Solaris 11 Package Repositories* manual at http://docs.oracle.com/cd/E23824_01/html/E21803/index.html

Both NIM and IPS provide command line as well as graphical user interfaces to perform software management tasks. The IPS `pkg(1)` command and its associated subcommands (such as `pkg install`, `pkg uninstall`, and `pkg list`) offer functionality that is similar to the `nimconfig` and `nimdef` NIM tools. While IPS includes two interfaces for software management (one GUI and one command line interface), it does not provide a terminal user interface such as the one available in AIX.

- In IPS, the Package Manager is used to search, install, and remove individual packages or groups of packages. Initiated on the command line by `packagemanager(1)`, it also is used to add, remove, and modify package publishers, or to create, remove, and manage Boot Environments. (Boot Environments are clones of the active boot image.)
- In IPS, the Update Manager, initiated on the command line by `pm-updatemanager(1)`, is a related GUI used to update all packages in an installation image for which updates are available. The Update Manager is similar to managing interim fix packages. Package Manager provides descriptive details about each package (including the versions, time stamps, and descriptions to simplify identification), making software management an easy and intuitive process. Package Manager groups packages into categories to simplify the task of locating a specific package.

Administrators can use IPS packaging tools in a zones-enabled environment. (Oracle Solaris Zones are an operating system virtualization technology used to provide isolated and secure execution environments. Each Oracle Solaris system hosts a global zone in which non-global zones can be created.) With few exceptions[3], running IPS commands in the global zone only impacts packages in the global zone. For example, executing `pkg install` in the global zone installs the package there; it is not propagated to any other zones. This model allows zones to be independently administered and maintained with their own separate software stacks. A zone administrator can use the same IPS packaging tools to manage software within a non-global zone.

See the *Oracle Solaris Administration: Oracle Solaris Zones, Oracle Solaris 10 Zones, and Resource Management* manual at http://docs.oracle.com/cd/E23824_01/html/821-1460/toc.html for more information.

## Updating System and Application Software

Patching has traditionally been the means by which software vendors have addressed security issues, bug fixes, performance improvements, and new features. For both earlier AIX and Oracle Solaris versions, patching could sometimes be a complex process that required detailed manual analysis to understand dependencies before applying the required patches. AIX addressed this complexity with tools such as (SUMA and NIM, generic fileset management utilities, and optionally Tivoli Provisioning Manager. In contrast, Oracle Solaris 11 eliminates the software maintenance model of patch analysis and application. Instead, the IPS software packaging model releases updated software packages that are already integration-tested before they are made available for download and installation, reducing the risk of incompatibilities between production software modules or problems resulting from human error. In addition, package contents that have not changed are not downloaded, resulting in faster updates that minimize network bandwidth consumption.

Update Manager in Oracle Solaris 11 reconciles all installed system software packages, updating packages as needed to bring the base operating system environment to a defined and integration-tested level. Both Update Manager and Package Manager check package versions against packages in a specified repository and identify available updates for installed packages. The "Update All" function in Update Manager updates all installed packages, as does "Update All" in Package Manager. Package Manager, however, allows administrators to add, remove, or update individual packages in between full system update operations. There is also a command line equivalent (`pkg update`) for updating all packages. If the `pkg update` command (with no Fault Management Resource Indicators [FMRIs] specified) is run in the system's global zone, it updates all *dependent system software packages* in both the global zone and any non-global zones. This keeps the zones synchronized.

For more information, see *Installing Oracle Solaris 11 Systems* at http://docs.oracle.com/cd/E23824_01/html/E21798/index.html and *Adding and Updating Oracle Solaris 11 Software Packages* at http://docs.oracle.com/cd/E23824_01/html/E21802/index.html.

[3] Key exceptions include entire operating system updates and patches that are designed for system-wide propagation.

The Tivoli Provisioning Manager is used in many AIX deployments. While Tivoli Provisioning Manager has been available for Oracle Solaris for many years, customers who have committed to using this software should consult with their Tivoli support contacts regarding the time frame for Oracle Solaris 11 support.

## Upgrading Software and Boot Environments

In AIX environments, system administrators use Fix Central, SUMA, NIM, and SMIT to load entire Service Packs (SPs). System backups are performed prior to updating a system. This remains IBM's recommended best practice for native AIX system administration.[4] To minimize downtime, an AIX administrator can create an `alt_disk_install` (by using **`mksysb`** to create an image on a separate disk or by cloning the existing system and applying the updates to it. One can then use the **`bootlist`** command to return to the previous level). This practice is in stark contrast to the new Boot Environments available in Oracle Solaris 11.

In Oracle Solaris 11, Boot Environments are the default and recommended best practice. Creating a Boot Environment (BE) takes advantage of the underlying Oracle Solaris ZFS file system technology built into the operating system. Fast snapshot and cloning capabilities are used to replicate the active operating system image. Because Oracle Solaris ZFS uses a copy-on-write technique, cloning a file system takes seconds, even for large disk arrays. The ability to clone the system eliminates the need to manually backup the system prior to creating a new BE. Because the old BE remains in place until deleted, several "generations" of configuration are readily available on short notice.

By default, a new BE is created automatically when certain system packages (key drivers and kernel components) are updated or when the administrator updates all packages (via an "Update All" in the GUIs or with `pkg update`). In these cases, Oracle Solaris 11 first clones the current BE, applies packaging changes to the clone, and automatically activates the updated BE so that it takes effect after a reboot. If a problem occurs, the administrator can easily roll back to the previous BE image. In this way, Oracle Solaris 11 provides an administrative safety net for upgrades and software changes, helping to improve availability. Because fast reboot is configured as the default, systems can switch to a new BE quickly, often within a few seconds.

Administrators can use the Oracle Solaris 11 `beadm(1M)` utility to manage BEs. The Package Manager GUI also supports the most common BE management tasks.

For more information, see the *Creating and Administering Oracle Solaris 11 Boot Environments* manual at http://docs.oracle.com/cd/E23824_01/html/E21801/index.html.

[4] See http://www14.software.ibm.com/webapp/set2/sas/f/best/Managing_AIX_Updates_Best_Practices.pdf

## Automating Installations of Multiple Enterprise Systems

The AIX NIM utility uses a client-server model to automate the installation of multiple systems across an enterprise. In Oracle Solaris 11, the Automated Installer (AI) provides equivalent functionality. AI automates and batches installations based on customized, standardized system profiles, enabling hands-free installation of multiple systems in large-scale environments. Similar to NIM, AI uses a client-server model. AI also leverages other Oracle Solaris 11 technologies, specifically the IPS packaging model, the Service Management Facility (SMF), and WAN-compatible networking protocols such as the Dynamic Host Configuration Protocol (DHCP), Preboot Execution Environment and Trivial File Transfer Protocol (PXE/TFTP), Hyper Text Transfer Protocol (HTTP), and the Domain Name Service (DNS) and Multicast Domain Name Service (mDNS) protocols, to provide operational flexibility.

The AI installation server houses a SPARC and x86 network boot image, installation instructions (called *AI manifests*), and optional system configuration (SC) profiles. Clients can be customized with installation parameters such as disk layout and software selection, and with system configuration parameters such as host name, network configuration, and user accounts.

An AI client first boots over the network, obtaining its network configuration and the location of the installation server via DHCP. Next, the client is configured and installed according to the AI manifest that matches the client's characteristics. AI installs a minimal network boot image on the client. The client subsequently completes the installation by accessing the IPS software repositories specified in the manifest. To customize the system after AI installation, the client applies an SC profile that configures the system using SMF services during the first boot process. In addition, virtualized environments are provisioned automatically, with non-global zones configured and installed during the first boot process after the AI installation completes.

For more information on the Automated Installer, For more information, see *Installing Oracle Solaris 11 Systems* at http://docs.oracle.com/cd/E23824_01/html/E21798/index.html.

### Building a Customized Distribution Image

In AIX it is often a good practice to create an `lpp_source` resource that contains all of the fileset updates on a NIM master. This resource contains the software images necessary for AIX installation. In Oracle Solaris, the Distribution Constructor command line utility is used to customize and build installation images. Checkpoints are performed during the construction process using Oracle Solaris ZFS. As a result, portions of the process can be restarted without going back to the beginning each time a change is made.

Distribution Constructor builds an image based on parameters specified in an XML manifest file. Sample manifests define preset, default values for an image and can be edited to further customize the resulting image. Distribution Constructor supplies sample manifests that build customized images similar to the Oracle Solaris Live Media for x86 image, ISO images for an Oracle Solaris 11 text installation, or ISO images for the Automated Installer. It builds an ISO image or a USB image that is based on a generated ISO image. However, a USB image can be used only on x86 systems.

For more information on the Distribution Constructor, see the *Creating a Custom Oracle Solaris 11 Installation Image* guide at http://docs.oracle.com/cd/E23824_01/html/E21800/index.html.
Note that Oracle cryptographically signs distributed packages. Local IT staff can sign their approved packages as well, preventing rogue administrators (outside of those trusted with signing keys) from installing unauthorized software.

# Chapter 3. Managing Data

AIX 7.1 and Oracle Solaris support many of the same file systems. Oracle Solaris includes additional file systems that offer new capabilities to AIX users. Moving data from AIX to Oracle Solaris cleanly and efficiently requires some knowledge of the target file system(s) so that IT staff can determine how best to move data from one platform to another. This chapter describes the supported disk-based, network-based, and virtual file systems in AIX 7.1 and Oracle Solaris 11. Other storage-related topics, including volume managers, data backup and recovery, shadow migrations, swap space, and data transformations also are discussed.

Table 3-1 lists the disk-based file systems supported on AIX 7.1 and Oracle Solaris. Oracle Solaris supports many of the same file systems as AIX 7.1, enabling users to simply mount existing file systems rather than migrate them.

**TABLE 3-1. AVAILABLE FILE SYSTEMS AND**

| FILE SYSTEM | DESCRIPTION | AIX | ORACLE SOLARIS |
|---|---|---|---|
| JFS and JFS2 | Journaled File System (IBM proprietary) | √ | — |
| Oracle Solaris ZFS | A general-purpose, enterprise-class file system that integrates traditional file system functionality with built-in volume management techniques and data services such as deduplication and compression | — | √ (Default file system) |
| CDRFS | Allows access to the contents of a CD-ROM | √ | √ |
| UDFS | Allows access to a DVD (read-only) | √ | √ |
| UFS | Traditional UNIX file system | — | √ Available for non-bootable file systems |

## AIX File Systems

While AIX supports a number of file systems, the primary disk-based file systems used are the IBM proprietary journaled file system (JFS) and enhanced journaled file system (JFS2). JFS was designed for high-throughput environments, and JFS2 builds upon JFS with support for large files. Both JFS and JFS2 provide smaller maximum file and file system sizes than Oracle Solaris ZFS. For example, JFS limits files to 64 GB and file systems to 1 TB, while JFS2 limits files and file systems to 4 PB.[5]

[5] Source: http://www.ibm.com/developerworks/aix/library/au-aix7optimize3/index.html

JFS2 supports additional enhancements, including encrypted file systems (EFS) and internal and external snapshot capabilities.

## Oracle Solaris ZFS

Oracle Solaris ZFS is a 128-bit file system that provides the scalability to store and manage virtually unlimited amounts of data. Individual files and file systems can scale to 16 exabytes (16 EB). In Oracle Solaris 11, Oracle Solaris ZFS is the default file system, and always is used as the root file system. (Other file systems are supported as data file systems.) Using Oracle Solaris ZFS as the root file system enables fast root file system snapshots and easy roll back to previous states. Read-only root file systems are permitted, enabling the environment to be locked down for added security. In addition, Oracle Solaris ZFS provides encryption, compression, and deduplication as file system properties, enabling any combination of these features to be used in a single pool.

Since AIX does not support Oracle Solaris ZFS, and Oracle Solaris does not support JFS, a data migration must be performed. Migrating data from JFS to an Oracle Solaris ZFS file system can be accomplished by:

- **Connecting the systems and using NFS export/import capabilities.** Typically it is most effective to employ the Oracle Solaris ZFS shadow migration feature[6],. The process involves making the source AIX file system read-only and available over NFS. On the Oracle Solaris system, an Oracle Solaris ZFS file system is created. Using Oracle Solaris ZFS shadow migration facilities, the directory structure of the source AIX file system is moved to the Oracle Solaris system, and data is copied in the background. The advantage of this procedure is that the new Oracle Solaris ZFS file system can be put to use immediately. If users try to access a file that has not yet migrated to the Oracle Solaris system, the file is migrated immediately and made available. The `shadowstat` command can be used to monitor progress. AIX encrypted file systems (EFS) can be migrated in the same fashion, with some care taken to ensure that the Oracle Solaris ZFS target is defined with encryption enabled. Because data is sent through the network in the clear, a secure subnet should be employed.
- **Using backup and restore utilities.** Administrators can backup the JFS file system using tools that are available on both systems and restore the data to an Oracle Solaris ZFS file system.

See the *How to Manage ZFS Data Encryption* article at http://www.oracle.com/technetwork/articles/servers-storage-admin/manage-zfs-encryption-1715034.html and for complete information on how to set up and administer Oracle Solaris ZFS file systems, see *Oracle Solaris Administration: ZFS File Systems* manual at http://docs.oracle.com/cd/E23824_01/html/821-1448/index.html

[6] http://docs.oracle.com/cd/E23824_01/html/821-1448/gkkud.html for details of shadow migration

## Volume Management

In AIX, a separate Logical Volume Manager (LVM) is needed to group storage into logical volumes. Oracle Solaris ZFS eliminates the need for separate volume management altogether[7]. Instead of creating virtualized data volumes, Oracle Solaris ZFS aggregates devices into a storage pool. The storage pool describes the physical characteristics of the storage and acts as an arbitrary data store from which file systems can be created. File systems are no longer constrained to individual devices.

Space from a single storage pool is shared dynamically between multiple file systems and parceled out as file systems request it. Physical storage can be added to storage pools dynamically, without interrupting services. When capacity no longer is required by a file system in the pool, it is made available to other file systems. Oracle Solaris ZFS supports storage pool sizes up to $3x10^{23}$ PB, and up to $2^{64}$ devices per pool.

## Redundancy

In AIX, LVM is used to create RAID volumes for greater data integrity and availability. Oracle Solaris ZFS provides two types of redundant configurations: mirrored pools and RAID-Z pools. RAID-Z configurations include RAID-Z (distributed parity), RAID-Z2 (distributed double parity), and RAID-Z3 (distributed triple parity). Oracle Solaris ZFS dynamically stripes data across all non-redundant, mirrored, and RAID-Z configurations. Redundant storage pools can be created easily with a single command.

RAID-Z uses parity, striping, and atomic operations to ensure automatic reconstruction of corrupted data. All data is protected by 256-bit checksums, and a self-healing feature automatically repairs corrupt data. Because the file system is always consistent, time-consuming recovery procedures such as `fsck` are not required even if the system is shut down in an unclean manner. Complex storage administration is automatic and simplified, reducing administrative overhead. For example, a redundant file system spanning multiple disks can be created with a single command. In addition, Oracle Solaris ZFS file systems are mounted automatically when created and remounted automatically when systems are rebooted.

Data migration from an AIX LVM is similar to the JFS and JFS2 migration process described earlier in this document.

## Hybrid Storage Pools

Oracle Solaris ZFS provides the ability to optimize data placement for fast access with Hybrid Storage Pools. Flash technology can be placed in a new storage tier to assist hard disk drives by holding frequently accessed data to minimize the impact of disk latencies and improve application performance. By using Flash devices to handle certain types of I/O, and hard disk drives to store massive data sets, a Hybrid Storage Pool delivers significant application performance gains without sacrificing capacity (Figure 3-1).

[7] Oracle Solaris supports the legacy Solaris Volume Manager (SVM) product, although systems cannot boot from an SVM root device in Oracle Solaris 11. While SVM is supported, using Oracle Solaris ZFS is generally a better alternative as it decreases administration overhead and provides increased functionality.

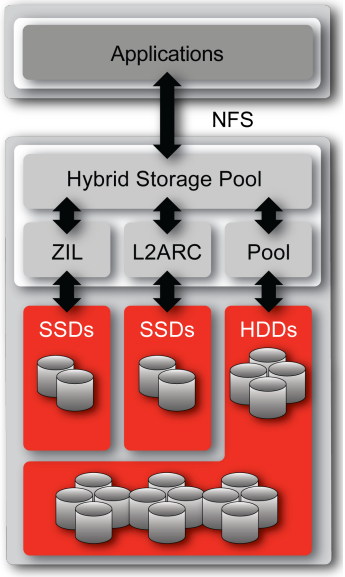


Figure 3-1. Hybrid Storage Pools optimize data placement to improve I/O performance.

Hybrid Storage Pool technology is designed to exceed the performance of Fibre Channel technologies without the additional management complexity and cost of a dedicated SAN. Several Oracle Solaris ZFS components are key to Hybrid Storage Pool operation and help accelerate application performance. The Oracle Solaris ZFS Adaptive Replacement Cache (ARC) is the main file system memory cache and is stored in DRAM. The Level Two Adaptive Replacement Cache (L2ARC) extends the ARC into read-optimized Flash devices to provide a large read cache to accelerate reads. The Oracle Solaris ZFS Intent Log (ZIL) is transactional and uses write-based Flash devices to provide a large cache to accelerate writes.

Sophisticated file system algorithms in Oracle Solaris ZFS use the ARC in memory and the L2ARC on Flash devices to determine pre-fetch or data placement during sustained read operations. Flash devices accelerate write throughput for Oracle Solaris ZFS synchronous write I/O operations, helping to boost write performance.

## Virtual File Systems

Table 3-2 lists common virtual file systems and details their availability on AIX and Oracle Solaris.

**TABLE 3-2. VIRTUAL FILE SYSTEMS**

| FILE SYSTEM | DESCRIPTION | AIX | ORACLE SOLARIS |
|---|---|---|---|
| CTFS | Contract file system, used to create, control, and observe contracts (primarily used by SMF) | — | √ |
| FDFS | File Descriptor File Systems, provides explicit names for opening files using file descriptors | — | √ |
| FIFOFS | First-in, first out file system, provides named pipe files that give processes common access to data | √ | √ |

**TABLE 3-2. VIRTUAL FILE SYSTEMS**

| FILE SYSTEM | DESCRIPTION | AIX | ORACLE SOLARIS |
|---|---|---|---|
| LOFS | Loopback file system, allows the creation of a virtual file system so that files can be accessed using an alternative path name (loopback device in AIX terminology) | √ | √ |
| MNTFS | Provides read-only access to the table of mounted file systems for the local system | √ | √ |

## Network-Based File Systems

Both AIX 7.1 and Oracle Solaris provide support for the Network File System (NFS) and Common Internet File System (CIFS) network-based file systems (Table 3-3). The NFS distributed file access protocol (versions 2, 3, and 4) is supported in both environments.

**TABLE 3-3. NETWORK-BASED FILE SYSTEMS**

| FILE SYSTEM | DESCRIPTION | AIX | ORACLE SOLARIS |
|---|---|---|---|
| NFS | Network File System | √ | √ |
| SMB (CIFS) | Server Message Block (SMB) service provides distributed resources to Windows and Mac OS systems and supports Common Internet File System (CIFS) | √ | √<br>Supports all Microsoft Windows access control list (ACL) semantics<br>Supports simultaneous export of NFS and CIFS |

The Server Message Block (SMB, also known as CIFS) service provides access to distributed files and directories to Microsoft Windows, Mac OS, and Oracle Solaris clients. The Oracle Solaris kernel provides a built-in Server Message Block (SMB) protocol server, and a client implementation supports numerous SMB dialects, including NT LM 0.12 and Common Internet File System (CIFS). While Oracle Solaris offers a fast, kernel-level CIFS implementation (available with the optional installation of the `service/file-system/smb` package), AIX provides CIFS support through the IBM `bos.cifs` fileset, which provides CIFS services via software running in user space.

See the "Managing Network File Systems" section of the *Oracle Solaris Administration: Network Services* guide at http://docs.oracle.com/cd/E23824_01/html/821-1454/rfsintro-2.html#scrolltoc for detailed information on how to set up and administer NFS file systems.
See the *Oracle Solaris Administration: SMB and Windows Interoperability* guide at http://docs.oracle.com/cd/E23824_01/html/821-1449/toc.html for detailed information on the SMB server and client in Oracle Solaris.

See the *Oracle Solaris Administration: Devices and File Systems* guide at http://docs.oracle.com/cd/E23824_01/html/821-1459/toc.html for information on memory-based file systems in Oracle Solaris.

## Swap Space

AIX implements paging and swapping very differently than other UNIX operating systems, including Oracle Solaris. While the fundamental configuration of physical storage as a backing store is similar, management is quite different. AIX memory management is based on the AIX Virtual Memory Manager (VMM)[8]. VMM uses swap space (paging) to store data that does not fit into RAM. A special logical volume is dedicated to VMM. Managing the space available to VMM is critical, as inadequate space can result in entire processes being lost or a system crash.

By default, Oracle Solaris reserves space for swap on the root storage pool (an Oracle Solaris ZFS file system). Swap devices are not pre-allocated to fixed-size slices, enabling the swap size to be modified at any time as needed. Additional swap volumes can be added to increase the amount of available swap space. In addition, system administrators can choose to allocate RAM as a high-speed memory store by creating a temporary file system (TMPFS) that serves as swap space. This gives system administrators greater manual control than typically is available under AIX.

For more information, see the *Oracle Solaris Administration: ZFS File Systems* guide at http://docs.oracle.com/cd/E23824_01/html/821-1448/index.html and *Adding or Changing Swap Space in an Oracle Solaris ZFS Root Environment* at http://docs.oracle.com/cd/E23824_01/html/821-1459/gizfl.html.

## Additional Storage Software

In addition to the built-in file system support described in the previous sections, the following data storage software and file systems are available on Oracle Solaris to fill specific storage and data requirements.

- Oracle's Sun Storage Archive Manager (SAM) software offers tiered data storage capabilities such as data classification, centralized meta-data management, policy-based data placement and migration. Oracle's SAM-FS is a self-protecting file system that offers continuous backup and fast recovery features.
- Oracle's Sun QFS software is a robust cluster file system intended for environments that share large data volumes. Sun QFS provides nearly raw device access to information and data consolidation for read/write file sharing. Similar to GPFS in AIX, the Sun QFS software often is used in conjunction with Oracle's Sun

[8] IBM also uses the VMM acronym in virtualized deployments with VMware software. In that context, *VMM* stands for Virtual Machine Manager.

Storage Archive Manager (SAM) to provide automatic policy-driven storage tiering and data lifecycle management.

More information on Sun QFS Software can be found at http://www.oracle.com/us/products/servers-storage/storage/storage-software/qfs-software/overview/index.html and the Sun SAM Software at http://www.oracle.com/technetwork/documentation/data-mgmt-software-194001.html.

## Data Backup and Restore

Both AIX and Oracle Solaris support a wide range of backup utilities, including those listed in Table 3-4.

**TABLE 3-4. BACKUP AND RESTORE UTILITIES**

| UTILITY | DESCRIPTION | AIX 7.1 | ORACLE SOLARIS |
|---|---|---|---|
| `cpio` | Saves and restores archives; copies files and directories while replicating the directory tree structure | √ | √ |
| `Backup, mksysb, sysback, splitvcopy` | AIX-specific backup utilities | √ | — |
| `pax` | Extracts, writes, and lists archive files; copies files and directories (newer version of `cpio` and `tar`) | √ | √ |
| `tar` | Extracts, writes, and lists archive files; copies files and directories (newer version of `cpio` and `tar`) | √ | √ |
| `zfs send zfs receive` | Send a snapshot to a file, file system, or system | — | √ |

Many backup utilities are common to both systems. For example, the newer `pax` utility is POSIX-compliant and is compatible with both AIX and Oracle Solaris. The **`pax`** utility supports a wide variety of archive formats, including **`tar`** and **`cpio`**, and can be used for data migration from one system to another. Data backed up using the **`pax`** utility on an AIX system can be imported on an Oracle Solaris system, and vice versa. Backup utilities also can be used to migrate data from one file system type to another on the same system.

There are a wide variety of third-party backup utilities that are available across platforms. Most enterprise production facilities employ one or more third-party backup application products. Popular third-party backup application tools run on both AIX and Oracle Solaris, such as Tivoli Storage Manager.

Experienced system administrators know that different platforms provide different access control list (ACL) facilities—and not all utilities treat ACLs appropriately, particularly when migrating data to new environments. Care should be taken to note the ACL information in effect on AIX and ensure it propagates with the same

properties on Oracle Solaris. In general, results should be consistent for NFS v4 to NFS v4 POSIX style ACLs when moving data from AIX to Oracle Solaris.

### Snapshots

In addition to backup utilities, Oracle Solaris ZFS includes *snapshot* capabilities—the ability to create a read-only copy of an Oracle Solaris ZFS file system or volume and restore it at a later time, if needed. Snapshots can be created almost instantly, and initially consume no additional disk space within the storage pool. Replicated streams of descendant file systems can be sent to named snapshots, preserving properties, snapshots, file systems, and clones. With snapshots, developers can save the state of a file system at a particular point in time, and recreate it on another machine to simplify data migration.

The JFS and JFS2 file systems in AIX provide two forms of snapshots: internal (within the file system) and external (between file systems the `backsnap` command). An AIX internal snapshot is equivalent to an Oracle Solaris ZFS snapshot, while the external snapshot capabilities is provided via the Oracle Solaris ZFS send and receive commands. For example, the `zfs send` command can be used to replicate a snapshot to a remote system, another file system on the same host, or to a "backup" file within the file system.

See the man pages at http://docs.oracle.com/cd/E23824_01/html/821-1461/index.html for detailed information on user commands and utilities, such as **`cpio`** and **`tar`**. For information on Oracle Solaris ZFS snapshots, see the "Overview of ZFS Snapshots" section of the *Oracle Solaris Administration: ZFS File Systems* guide at http://docs.oracle.com/cd/E23824_01/html/821-1448/gbciq.html.

## Data Transformation

*Data transformation* is the process of converting data from one format to another, and is an important component of any migration effort if data is to be readable on the target system. Data transformation can involve file systems, file content, applications, and database content.

### Encoded Data Transformations

Encoded data transformations are necessary when data is stored in a different or incompatible file format than the receiving system anticipates. Fortunately, AIX and Oracle Solaris typically use contemporary standardized text formats to store textual data, such as ISO, POSIX 1003.2, and X/Open (XPG3+). As a result, conversion issues are minimized.

Applications that depend on the AIX `libi18n.a` library, EBCDIC, or early IBM proprietary encodings such as IBM 281 for Japanese and IBM 500 for Swiss German, should be upgraded to use standardized encodings before the migration to Oracle Solaris.

The standard `dd` command can be used to do simple stateless EBCDIC to ASCII conversions, but a variety of national character set issues can cause complications. The Oracle Solaris `convert_external` procedure located in the `libm` library can be used to translate S/370 and various other floating-point data formats to the

IEEE Standard used by Oracle Solaris. Nonetheless, converting and validating the conversion prior to the migration generally is preferred.

For more information, see the *International Language Environment Guide* at http://docs.oracle.com/cd/E23824_01/html/E26033/index.html.

### Application Data Transformation

AIX and Oracle Solaris provide many common applications and utilities for managing data. For example, the tape archive utility (`tar`) uses a similar data format and provides many common options in both environments. This commonality is true for many other applications and utilities, and can yield significant benefits both during and after the data migration. For those applications that differ between AIX and Oracle Solaris, most provide a utility to convert standard data interchange formats, such as comma-separated values or tab delimited files, into their format.

### Database Transformation

Many enterprise applications depend on large databases. If an AIX environment uses an older version of a database, licenses may or may not be available for those versions on Oracle Solaris. IT organizations should be prepared to acquire a current version of the database software. It is important to note that changes to the infrastructure may be needed to support the new database and its configuration. Existing data should be immediately accessible.

Typically, DB2 and Oracle Database are used in AIX production environments. When the same database software is available in both environments, migrating to Oracle Solaris may be as simple as exporting the database running on AIX to a standardized file format, followed by an import into a new database on Oracle Solaris. Such migrations are relatively easy.

When the port also involves a change in database vendors, more extensive data transformations are typically required. Because database transformations are usually such a large part of the overall migration effort, many specialized utilities have been created to address them. These programs, called Extract, Transform, and Load (ETL) utilities, take a wide array of formats and convert them into Structured Query Language (SQL) for relational database management systems (RDBMS). Most RDBMSs provide a basic set of utilities to convert SQL or standard interchange formats into their data storage format.

# Chapter 4. Virtualizing Infrastructure

Access to robust virtualized environments is essential for consolidation efforts and the move to cloud computing. Both AIX and Oracle Solaris provide extensive and powerful virtualization functionality. While both environments provide a wide range of virtualization technologies, there is no simple one-to-one mapping. IT staff should carefully consider their application requirements, system resources, and service-level agreements and make adjustments to the virtualization deployment as part of a careful migration. Both AIX and Oracle Solaris provide facilities for partitioning and core sharing that are implemented as hardware or software partitions, hypervisor-based virtual machines, and operating system virtualization, as well as networking and I/O virtualization.

It is important to note that Oracle Solaris is a completely virtualized operating environment, with technologies that span server, network, and storage virtualization to help IT organizations optimize enterprise infrastructure resources. Because no two deployment environments have exactly the same needs, Oracle's virtualization technologies provide varying degrees of isolation, resource granularity, and flexibility, and can be used separately or (more often) together to tackle specific deployment environment challenges.

## Hardware Partitioning

Hardware partitioning in Oracle Solaris is provided by Dynamic Domains on Oracle's SPARC Enterprise M-Series servers. The maximum number of Dynamic Domains is model dependent, ranging from none on entry-level systems to 24 on the SPARC Enterprise M9000 server. These systems provide electrical isolation as well as logical isolation, and employ highly redundant hardware. This type of hardware isolation is available to AIX only on IBM zSeries hardware.

More information on Dynamic Domains can be found in the *Oracle SPARC M-Series Servers* documentation at http://www.oracle.com/technetwork/documentation/sparc-mseries-servers-252709.html?ssSourceSiteId=ocomen.

## Virtual Systems

Over time, IBM introduced a variety of terms related to its virtualization technologies. Today, IBM refers to the group of technologies—LPARs, DLPARs, and SPLPARs—under the PowerVM name. Most AIX administrators use the name LPARs when referring to any of these capabilities.

Oracle Solaris provides two fundamental mechanisms for server virtualization.

### Virtual Machines

IT organizations that use the Logical Partitions (LPARs) and Micro-Partitioning facilities of PowerVM on AIX systems can achieve similar levels of partitioning and isolation using Oracle VM Server for SPARC (previously called Sun Logical Domains) in conjunction with Oracle Solaris Zones. Purpose-built for Oracle's SPARC T-Series servers with chip multithreading technology, domains provide a full virtual machine that runs independent operating system instances and contains a wide range of virtualized devices. Like PowerVM,

domains use a hypervisor that resides within hardware at the chip level. Because the software is tightly integrated with the hardware, virtual machines can take advantage of underlying system advancements and have little overhead, unlike the pure software-based solutions provided by many other vendors. For example, Oracle VM Server for SPARC does not over subscribe virtual machines to a processor thread, eliminating the need for a scheduler in the on-board hypervisor. In fact, the SPARC T4 processor offers 64 threads, each capable of running a virtual machine. Furthermore, CPU loads are balanced by a logical domain manager running in the control domain. As a result, Oracle VM Server for SPARC does not require a workload manager to be embedded in the hypervisor, creating a more lightweight, faster, and reliable solution.

Generally speaking, the number of domains available in a SPARC T-Series system is the same as the number of hardware threads. The smallest domain is one processor thread, but more commonly at least a full core is dedicated to each domain. On AIX, Micro-Partitioning has historically scaled down to one-tenth of a core[9], and each Micro-Partition has its own AIX operating system instance.

The virtual machine facilities in both AIX and Oracle Solaris provide excellent scalability, fully virtualized resources, dynamic configuration, virtual I/O, and live migration[10].

### Operating System Virtualization

Workload Partitions (WPARs) can be mapped to Oracle Solaris Zones. While WPARs have been available to AIX customers since AIX 6.1, they have not generally been deployed in production settings. More recently, organizations migrating from AIX 5.x to AIX 7 have been encouraged by IBM to utilize WPARs in AIX 7. In contrast, Oracle Solaris Zones are a popular technology in development and production environments. Because they are built into the operating system, Oracle Solaris Zones can be used on Oracle's entire range of SPARC and x86 processor-based servers, enabling IT organizations to standardize on a virtualization technology across server architectures. Key capabilities include:

- **High consolidation ratios.** Oracle Solaris Zones can be used to provision many secure, isolated runtime environments for individual applications using flexible, software-defined boundaries. Because they run under a single operating system kernel, Oracle Solaris Zones enable fine-grained control over rights and resources within a consolidated server without increasing the number of operating system instances to manage. Computing resources—CPUs, physical memory, network bandwidth, and more—can be dedicated to a single application one moment and shared with others in an instant, without moving applications or rebooting the system, dynamic domain, or logical domain where the Oracle Solaris Zone resides. In addition, extremely low overhead enables IT administrators to achieve high consolidation ratios.
- **Simplified consolidation.** A preflight checker, `zonep2vchk(1M)`, can help system administrators identify identifies issues that could affect the migration from physical to virtual servers, and creates zone configuration output for the target zone.

[9] Current versions of AIX can optionally scale in 1/100th of a core steps.

[10] Oracle VM Server for SPARC performs encrypted live migration by leveraging on-chip cryptographic acceleration.

- **Legacy operating system support.** Oracle Solaris Zones provide a solid bridge for running older Oracle Solaris instances with minimal overhead, with the ability to create Oracle Solaris 10 environments on Oracle Solaris 11. Organizations can use this feature to reap the benefits of Oracle Solaris 11 without disrupting existing applications, run legacy Oracle Solaris 10 applications, or test them on Oracle Solaris 11 platforms[11].
- **Rapid application deployment.** With Oracle Solaris Zones, applications can be developed in an isolated environment and packaged for movement to testing systems. Shared storage makes it possible for the transition to happen quickly, and applications do not need to be duplicated. Once tested, applications can be moved quickly to production systems. With these capabilities, organizations can experience rapid roll out of applications, little downtime, and automatic roll back to development and testing systems when needed.
- **Rapid deployment and patching.** Oracle Solaris enable IT administrators to rapidly and easily deploy one or more pre-built, pre-configured, pre-patched virtual environments and their application stacks. In addition, the ability to update boot environments in parallel enables IT administrators to effectively patch virtual environments in parallel, with no downtime.
- **Shared storage.** In Solaris 11.1, it is easy to host an Oracle Solaris Zone on arbitrary storage device objects, including Fibre Channel or iSCSI targets, providing automatic zpool encapsulation. See `zonecfg(1M)` for details.
- **Immutable zones.** Immutable zones allow the creation of a read-only copy of a zone and related Oracle Solaris ZFS file system, including the boot and root file systems. The ability to deploy read-only zones provides an additional security barrier to lock down applications and data resident to that zone. Providing read-only access to an application and its data further secures it and prevents unauthorized access or hacking. In addition, Oracle Solaris 11 network services are disabled by default in an immutable zone, or set to listen only for local system communications, to limit opportunities for unapproved access. Unless performed as specific maintenance operations, modifications to system binaries or system configurations are blocked.

  Immutable zones can be a key part of a security-in-depth implementation. For example, system administrators can deploy immutable zones to protect the environment and data link protection to protect the network. In addition, Oracle Solaris ZFS encryption and read-only access can be used within the zone to protect data. More information about how these features can be used together to provide security-in-depth can be found at https://blogs.oracle.com/darren/entry/encrypted_immutable_zones_on_shared.

## Key Virtualization Similarities and Differences

In AIX, LPARs expanded over time to include virtual networking, Active Shared Memory, and Micro-Partitions, all of which permit selective sharing of LPAR resources. In contrast, the operating system virtualization capabilities in Oracle Solaris were designed from the ground up to maximize server resource

[11] This often is helpful for running third-party software that is not yet certified on Oracle Solaris 11.

sharing and utilization. While some AIX environments that use LPARs with Shared Active Memory and Micro-Partitions map well to domains, others are better served by transitioning to Oracle Solaris Zones. AIX LPARs, Oracle Solaris Zones, and domains provide:

- **Isolation.** Both AIX LPARS and Oracle VM Server domains provide complete operating system isolation. Oracle Solaris Zones offer comprehensive isolation, particularly in the case of immutable zones, but with greater resource sharing than early AIX LPARs.
- **Scalability.** PowerVM Enterprise edition provides support for up to 254 LPARs per server. These LPARs can be subdivided further using Micro-Partitions. Oracle VM Server for SPARC supports many domains, with the precise number supported on a given platform dictated by the system's configuration. In addition, domains use logical domain channels to provide a conduit for services, such as networks and shared devices. The number of channels depends on the server, ranging from 512 on entry-level SPARC T-Series servers to as many as 768 on high-end SPARC T4 servers. See the "Oracle VM Server for SPARC 2.2 Release Notes" at http://docs.oracle.com/cd/E35434_01/html/E23810/knownissues.html for details.
- **Workloads**. Just as different IBM LPARs can run different AIX versions, different domains can run different Oracle Solaris versions. In addition, Oracle Solaris Zones can support older versions of the operating system. System administrators typically mix domain and zone virtualization features to gain the best hardware utilization while preserving application and user expectations.
- **Security.** IBM LPARs, Oracle Solaris Zones, and domains all provide extensive security features. While there are differences in the details, IT organizations can be confident that a migration from AIX to Oracle Solaris is not going to reduce the security of virtualized systems. See the security chapter later in this document for more details.
- **Live migration.** An active domain can be migrated to another physical machine, even across different generations of SPARC T-Series servers. Domains are encrypted automatically during migration at near wire speed by leveraging the on-chip cryptographic acceleration in Oracle server hardware.

For more information on Oracle Solaris Zones, see *Oracle Solaris Administration: Oracle Solaris Zones, Oracle Solaris 10 Zones, and Resource Management* at http://docs.oracle.com/cd/E23824_01/html/821-1460/index.html.
For more information about Oracle VM for SPARC, see the Oracle VM Server for SPARC documentation at http://www.oracle.com/technetwork/documentation/vm-sparc-194287.html.

## Network Virtualization

Unlike the hypervisor approach used in AIX environments, which limits network virtualization to network interfaces, Oracle Solaris virtualizes all aspects of network topology within a server virtualization framework. Virtualization and resource control reside directly in the operating system to eliminate layered functionality with heavy overhead and undue complexity.

- **Aggregate and share network resources.** The network virtualization capabilities built into Oracle Solaris enable organizations to take advantage of all available bandwidth and create scalable network infrastructure. Using the basic building blocks of VNICs, virtual switches, virtual interconnects, virtual LANs, virtual routing and virtual firewalls, high-bandwidth physical network connections can be carved up to enhance network utilization or aggregated as needed to meet peak workload demands. Similarly, multiple Gigabit Ethernet connections can be aggregated to offer a single, larger network connection with greater bandwidth to applications. Fine-grained monitoring and control facilities ensure systems and applications have access to needed bandwidth.
- **Consolidate workloads and optimize network utilization.** Built-in network virtualization promotes more effective sharing of network resources and enhances the ability to consolidate server workloads. In Oracle Solaris, network traffic can be isolated and assigned limits or guarantees on the amount of bandwidth it can use. Placing limits on bandwidth consumption improves network utilization and performance rates and supports operating system virtualization, cloud computing, and consolidation efforts.
- **Gain flexibility without sacrificing performance.** Network Virtualization virtualizes the network stack and network interface card (NIC) around any service protocol, such as HTTP, HTTPS, FTP and NFS, or virtual environments created with Oracle Solaris Containers or Oracle VM Server for SPARC. Each virtual stack can be assigned its own priority and bandwidth on a shared NIC without degrading performance. In addition, network workloads can be parallelized across multiple processors, cores, and threads to increase performance.

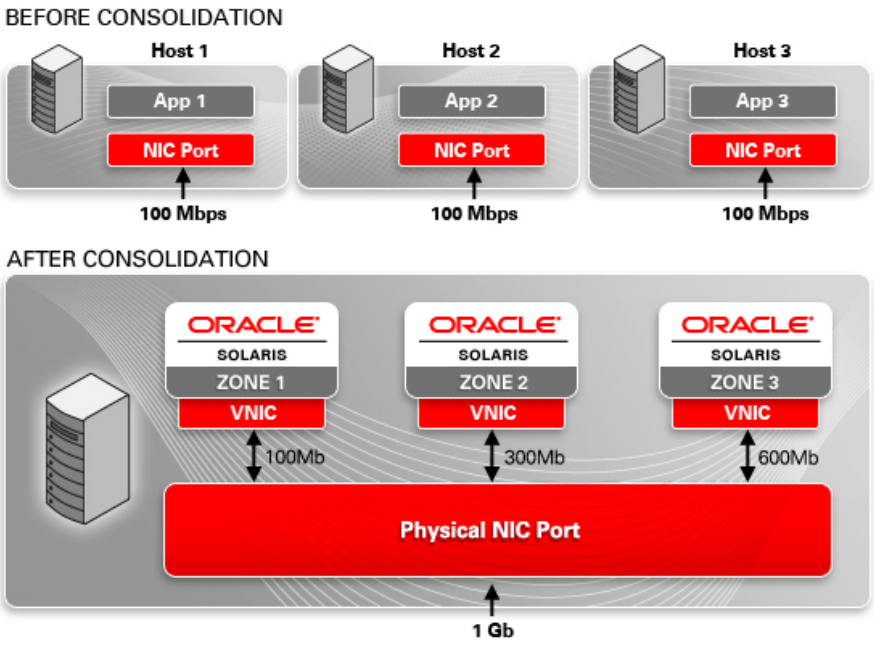


Figure 4-1. Networks can be virtualized to optimize resource sharing and bandwidth utilization.

- **Manage traffic.** Oracle Solaris support data center bridging. This technology enhances the ability of traditional Ethernet networks to manage traffic, especially in environments with high network traffic volume and transmission rates. While Fibre Channel can be dedicated to this type of traffic, using dedicated links to service only one type of traffic can be costly. As a result, Fibre Channel over Ethernet (FCoE) is more commonly used, with data center bridging addressing Fibre Channel's sensitivity to packet loss while traversing Ethernet networks. With DCB, peers can distinguish traffic based on priorities. As a result, hosts can ensure packet integrity is preserved for high-priority traffic in the event of congestion between hosts. Using the DCB exchange protocol (DCBX), communicating hosts can exchange configuration information that affects high-speed network traffic. Peers can negotiate on a common configuration that ensures continuous traffic flow while preventing packet loss for high-priority packets.
- **Model networking environments.** IT organizations can model a complete data center networking topology in a virtual network, reducing the limitations created by the availability of physical NICs.
- **Implement edge virtual bridging (EVB).** Part of an evolving IEEE standard, EVB allows a host to exchange virtual link information, such as bandwidth share and priority definitions for physical links, with external switches. The ability to enable reflective relay on the external bridge port can be used to ensure communication between virtual machines occurs without packets traversing the network. In addition, administrators can automate virtual port configuration on the bridge. For more information, see http://docs.oracle.com/cd/E26502_01/html/E28993/gmhiw.html.
- **Secure applications from attacks.** The architecture dynamically manages bandwidth resources, and can provide a better defense against denial-of-service attacks directed at a particular service or virtual machine by isolating the impact to that entity. As a result, administrators can ensure key applications are not starved of network resources, keep critical backup operations from impacting network performance, and control the resources users receive more effectively. While bandwidth management does not prevent denial-of-service attacks, it does ensure a compromised zone cannot become a perpetrator of denial-of-service attacks against other systems.

See *Oracle Solaris System Administration: Network Interfaces and Network Virtualization* at http://docs.oracle.com/cd/E23824_01/html/821-1458/index.html for more information on Network Virtualization. See http://docs.oracle.com/cd/E26502_01/html/E28993/gmfbf.html for more information on Data Center Bridging.

## I/O Domains

While named differently in the two environments, both AIX and Oracle Solaris support the notion of an I/O domain. In Oracle Solaris, an I/O domain has direct ownership of, and direct access to, physical I/O devices.

- **High performance.** As IT organizations consolidate and virtualize more servers, software emulated I/O is exposed as a limiting factor. Similar to AIX, Oracle Solaris 11 supports the single-root I/O virtualization (SR-IOV) framework, defining extensions to the PCI Express (PCIe) specification to support the efficient

sharing of PCIe devices among virtual machines in hardware and software. Using Oracle VM Server for SPARC, IT organizations can create special-purpose I/O domains that assign PCIe bus, network interface units (NIUs), or PCIe endpoint devices to a domain. Because they have direct access to I/O devices, I/O domains enable more I/O bandwidth to be available for services that support high-performance applications, as well as virtual I/O services for guest domains.

- **High availability.** Multiple I/O domains can be used to eliminate single points of failure. For example, guest domains can be configured with virtual disks and virtual devices that are provisioned from multiple service domains. In such a configuration, the failure of a service domain or I/O path or device does not result in an application outage.
- **High uptime with rolling upgrades.** The ability to provision virtual infrastructure from multiple devices supports IT efforts to maximize uptime even during the upgrade process. Domains can be upgraded individually, while their guests continue to operate. While a given device or path is involved in the upgrade process, other devices and paths continue to provide services to guests.

## Summary of Comparable Virtualization Technologies

Table 4-1 summarizes the key differences between AIX and Oracle Solaris virtualization technologies.

**TABLE 4-1. COMPARISON SUMMARY OF AIX AND ORACLE SOLARIS VIRTUALIZATION TECHNOLOGIES**

| VIRTUALIZATION TECHNOLOGIES | | |
|---|---|---|
| **TYPE** | **AIX** | **ORACLE SOLARIS** |
| Hardware Partitions[12] | LPAR (zSeries) | • Dynamic Domains (Oracle's SPARC Enterprise M-Series servers)<br>• Up to 24 domains per system |
| Virtual Machines | LPARS | • Oracle VM Server for SPARC (Oracle's SPARC T-Series servers)<br>• Oracle VM VirtualBox (x86 servers only) |
| Operating System Virtualization | WPARs | • Oracle Solaris Zones (Oracle's SPARC and x86 servers) |
| Network Virtualization | Built into the Operating System (VIOS) | • Built into the operating system<br>• SR-IOV framework |
| Memory Sharing | Active Memory Sharing | • Oracle Solaris Zones |

[12] The term *hard partition* has multiple meanings, depending on the context. For the purposes of Table 4-1, the intention is complete hardware separation so we are using the term hardware partition. IBM zSeries (and EC12) and Oracle's SPARC Enterprise M-Series servers provide this level of hardware partitioning. Some licensing agreements, including some Oracle licenses, use vendor facilities such as Oracle Solaris capped containers as "hard" partitions. See http://www.oracle.com/us/corporate/pricing/partitioning-070609.pdf and http://www.oracle.com/technetwork/server-storage/vm/ovm-sparc-hard-partitioning-1403135.pdf for additional information.

# Chapter 5. Keeping Systems, Applications, and Services Available

Understanding the need to deliver reliable infrastructure, AIX and Oracle Solaris both include tools for monitoring and managing systems and processes, as well as handling automatic failover, restart, and recovery in the event of a disruption in service due to hardware or software malfunction. Table 5-1 maps technologies for maximizing availability in AIX 7.1 deployments to counterpart technologies for deploying highly available Oracle Solaris systems. While AIX 7.1 is cluster aware, an additional tool (PowerHA) is required to provide failover protection. Oracle Solaris 11 is cluster aware, with Oracle Solaris Cluster delivering cluster deployment.

**TABLE 5-1. AVAILABILITY TECHNOLOGY MAPPINGS**

| FUNCTIONALITY | AIX | ORACLE SOLARIS |
|---|---|---|
| Fault detection, reporting, and recovery | FFDC (First Failure Data Capture) | Fault Management Architecture (FMA) |
| Configuring and managing system services | `SMIT` | Service Management Framework (SMF) |
| Clustering | IBM Power HA | Oracle Solaris Cluster |
| Infrastructure and system monitoring and management | Cluster Aware AIX (CAA), RSCT and Power HA | Oracle Enterprise Manager and Oracle Enterprise Manager Ops Center |
| DTrace | | DTrace facilitates rapid identification of hung processes, bottlenecks, and other performance issues that impact system availability |

## Predictive Self Healing

Predictive Self Healing technologies proactively monitor and manage system components to help organizations optimize IT service availability. These technologies are built into Oracle Solaris to leverage hardware diagnostics, allowing business-critical applications and essential system services to continue uninterrupted in the event of software defects, major hardware component failures, and even misconfigured software. The Oracle Solaris Service Management Facility (SMF) and Oracle Solaris Fault Management Architecture (FMA) are the two main components of Predictive Self Healing.

### Oracle Solaris Service Management Facility

AIX provides SMIT for configuring services and system daemons. It supports a variation of the traditional UNIX run levels (called run levels a–e), and typically only the customary run level 2 is used. Oracle Solaris includes a standardized infrastructure for controlling system services—the Service Management Facility (SMF)—that augments the traditional `/etc/rc` start and stop scripts, `init` run levels, and configuration files.

SMF provides a framework that simplifies the management of system services and delivers improved ways to manage them. Services are treated as objects that administrators can consistently configure, enable, control, observe, and manage in a uniform way. Relationships and dependencies between services are easily defined and managed—an advantage over traditional UNIX `/etc/rc` scripts. Information needed to manage each service is stored in a service repository. In the event of a failure, services are restarted automatically (along with any dependent services) and reports are generated that identify the failing component and services that depend on the failing component. This is done whether the service was accidentally terminated by an administrator, aborted as the result of a software programming error, or interrupted by an underlying hardware problem. Administrators can be notified of service state transitions and fault management events via SNMP traps or email messages, providing better visibility into errors and improving debugging capabilities to help resolve service-related problems quickly. Application developers can add the appropriate logic to their installation packages, enabling layered services to be integrated into the SMF framework and managed by native Oracle Solaris facilities.

Along with the installation and packaging technologies in Oracle Solaris, SMF is at the heart of initial system configuration tasks and a key part of the underlying software installation architecture. Different SMF services are activated on first reboot as a part of the operating system installation process, applying system configuration profiles to configure and activate services. During software package installations on a running system, SMF services can apply or refresh configuration caches as an alternative to post-installation scripts, at the same time taking into account defined service dependencies. SMF helps to apply configuration changes in a reliable and repeatable fashion, enabling more seamless and error-free software installations and upgrades.

### Fault Management Diagnosis and Recovery

To increase system availability, the Fault Management Architecture (FMA) in Oracle Solaris helps to detect system problems, similar to the way System Fault Management (SFM) does in AIX. FMA goes further than simple detection and reporting by diagnosing faults and initiating recovery measures that can help to prevent outages. FMA tries to configure problem components out of a system before a failure occurs. In the event of a failure, it initiates automatic recovery using SMF. FMA builds a suspect list of root causes based on error patterns and identifies the likely associated system resources. Following this diagnosis step, FMA provides fault information to agents that know how to respond to specific faults.

At a high level, the FMA stack contains error detectors, diagnosis engines, and response agents. A Fault Management daemon (itself a service under SMF control) connects FMA components and acts as a multiplexor between them. FMA error detectors sense errors in the system and report them to a diagnosis engine. The diagnosis engine interprets the report, determines whether a fault or defect is present, and identifies a probable cause. The source of the problem may have an associated Automatic System Reconfiguration Unit (ASRU) or a Field-Replaceable Unit (FRU). An ASRU is a system resource that an FMA agent can disable to isolate the problem and suppress further error reports. In many cases an FRU can immediately be removed from the service, mitigating the problem until replacement is possible.

Oracle Solaris notifies administrators of FMA fault management events as well as SMF service state changes. Administrators can configure Simple Network Management Protocol (SNMP) trap notifications and Simple Mail Transport Protocol (SMTP) email notifications to watch for certain events or services. Systems often can automatically configure around a failed component and notify the system administrator of such an event.

Notifications also can be sent directly to Oracle with Automated Service Requests (ASR), providing automatic telemetry for customers who have active Oracle support agreements and enabling a proactive response from Oracle service and support engineers.

For more information on SMF and FMA see *Oracle Solaris Administration: Common Tasks* at http://docs.oracle.com/cd/E23824_01/html/821-1451/index.html.

## Clustering Technology

Strategic business applications demand continuous availability, so IT departments often configure clustering technologies when deploying data center systems to meet strict service-level agreements. Because organizations can implement geographically dispersed clusters, clustering technologies also are implemented to meet requirements for disaster recovery. For organizations that deploy local, campus, metropolitan, or worldwide clusters, Oracle Solaris Cluster extends the high availability features built into the core Oracle Solaris operating system.

Oracle Solaris Cluster provides clustering functionality similar to the combination of IBM's Cluster Aware AIX (CAA), Reliable Scalable Cluster Technology (RSCT) (3.1 and later), PowerHA (7.1 and later), VIOS, and third-party ISVs, service providers, and software products. At a high level, AIX cluster robustness relies on a combination of SAN and Ethernet networking, a special "gossip" protocol, and close cooperation of these various software layers.

In contrast, Oracle Solaris Cluster is a comprehensive clustering framework that encompasses the core Oracle Solaris Cluster software, Oracle Solaris Cluster Geographic Edition, Oracle Solaris Cluster agents, and developer tools and support for clustering commercial and open-source applications. To optimize the availability of mission-critical applications and services in traditional or virtualized environments, Oracle Solaris Cluster provides load balancing, automatic fault detection, and failover.

### Overview of Oracle Solaris Cluster

At its simplest, Oracle Solaris Cluster monitors the health of cluster components, including the stack of applications, middleware, operating system, servers, storage, network interconnects, and even Oracle Solaris Zones. Any failure executes a policy-based, application-specific recovery action. Recovery is enabled through redundant infrastructure and intelligent software algorithms.

From a physical perspective, an Oracle Solaris Cluster system consists of two or more nodes that work together as a single entity to cooperatively provide applications, system resources, and data to users (Figure 5-1). Each node provides some level of redundancy. Data is stored on highly available redundant disk systems, which may be mirrored, supporting data access in the event of a service interruption on a single disk or storage subsystem. Redundant connections are provided to disk systems so that data is not isolated in the event of a server, controller, or cable failure. A high-speed, redundant, private interconnect provides access to resources across the server set. Redundant connections to the public network also provide each node with

multiple paths for access to outside systems, helping ensure continued access in the event of a network connection or node failure.

No single hardware, software, storage, or network failure can cause the cluster to fail. Loss of service is prevented through hardware redundancy, hardware and software failure detection, automatic recovery of services, and application failover. In addition, a single management view enables the entire cluster to be managed as a single entity, reducing the risk of errors.

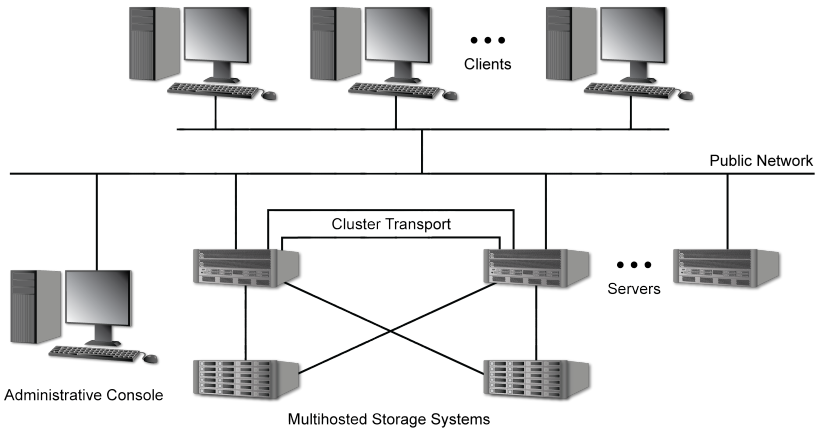


Figure 5-1. Oracle Solaris Cluster enables multiple servers and storage systems to act as a single system.

Oracle Solaris Cluster includes capabilities to detect, isolate, and contain failing cluster nodes. It accomplishes this using a robust, kernel-based membership monitor. Each node in the cluster sends out low-level Data Link Provider Interface (DLPI) packets once per second (a heartbeat) to each of its peers on each of the private networks. These packets are sent in the kernel interrupt context, making them very resilient to peaks in system load. A network, or path, between two nodes is declared down only if a heartbeat message does not complete the round trip between nodes over that specific path within the timeout period.

### Network Availability

Oracle Solaris Cluster leverages Oracle Solaris IP network multipathing (IPMP) as public network interfaces for monitoring local failures, and for performing automatic failover from one failed network adapter to another. IP network multipathing enables a server to have multiple network ports connected to the same subnet. First, IP network multipathing software provides resilience from network adapter failure by detecting the failure or repair of a network adapter. The software simultaneously switches the network address to and from the alternate adapter. When more than one network adapter is functional, IP network multipathing increases data throughput by spreading outbound packets across multiple adapters.

For scalable data services, requests go through a round-robin load-balancing scheme for a balanced load distribution to the various instances of the distributed application running within the cluster. Scalable data services can be made more secure through the use of Internet Protocol security (IPsec) services in combination with Oracle Solaris Cluster load balancing services.

### Data Integrity

Because cluster nodes share data and resources, Oracle Solaris Cluster works to ensure a cluster never splits into separate, active partitions that continue to access and modify data. Similar to PowerHA, Oracle Solaris Cluster applies fencing techniques and a quorum to protect data integrity. Failing nodes are isolated from the cluster and prevented from accessing clustered data. The fencing protocol can be chosen per storage device.

In a more complex situation where all paths across the private interconnect fail and the cluster breaks into multiple partitions, Oracle Solaris Cluster uses a quorum mechanism to recreate the cluster and resolve partitions or split brain syndrome, and to protect data integrity. The quorum also prevents amnesia by detecting and rejecting the use of outdated configuration information that could lead to data corruption. In PowerHA configurations, split brain syndrome is avoided by utilizing the special AIX "gossip" protocol.

The Oracle Solaris Cluster quorum can be tailored to the storage and system topology, enabling disk-based and software quorum solutions. A quorum device protocol permits the use of different types of disks, such as high-capacity 2 TB disk drives, SATA, and Flash as quorum devices. All quorum devices are continuously monitored to enhance availability.

Note: the Oracle Solaris `snoop_interval` variable set in the `/etc/system` file is the equivalent of the AIX `deadman` switch.

### Virtual Clustering

Oracle Solaris Cluster software also supports virtual clustering, allowing Oracle Solaris Zones to function in the role of cluster nodes. Virtual clusters allow organizations to deploy multiple applications or multi-tiered workloads on a single physical cluster configuration. Applications can run within a specific zone cluster under separate zone policy-based management. In the event of a zone failure, individual zones can be restarted or failed over. In this way, Oracle Solaris Cluster can protect applications that run in Oracle Solaris Zones, or Oracle Solaris 10 zones hosted on Oracle Solaris 11 systems.

### Key Components

Key components of Oracle Solaris Cluster include:

- **High availability framework.** The framework detects node failures quickly and activates resources on another node in the cluster. It includes a Cluster Membership Monitor, a distributed set of algorithms and agents that exchange messages over the cluster interconnect to enforce a consistent membership view, synchronize reconfiguration, handle cluster partitioning, and help maintain full connectivity among all cluster members. Inter-node message delivery and responses are handled in an atomic manner that accounts for delivery failures, node membership, and software revision level to provide for rolling upgrades.
- **Failover, scalable, and cluster-aware agents.** Failover and scalable agents are software programs that enable Oracle or third-party applications to take full advantage of Oracle Solaris Cluster features. Cluster-aware applications have direct knowledge of Oracle Solaris Cluster systems, such as Oracle Real Application Clusters (Oracle RAC) software. Oracle Solaris Cluster agents specify the actions to be taken should an application fail. Many agents for Oracle and third-party enterprise applications are available (see http://www.oracle.com/technetwork/server-storage/solaris-cluster/index.html). Oracle Solaris Cluster includes built-in support for Oracle Solaris 11 services such as Apache, Apache Tomcat, DHCP, DNS, NFS, as well as additional Oracle software such as Oracle WebLogic Server and Oracle Database (single

instance and Oracle RAC). If an agent does not exist for an application, one can be built using provided tools.

- **Highly available private interconnect.** Multiple types of interconnect technologies are supported by Oracle Solaris Cluster to establish a private communication channel between cluster nodes. Support for multiple interconnects helps to ensure high availability and improve performance of private inter-node communication. Heartbeats monitor cluster nodes over the private interconnect. If a server goes offline and ceases its heartbeat, it is isolated. Applications and data are failed over to another server quickly and transparently to users.
- **Live migration.** Domains created with Oracle VM Server for SPARC offer live migration facilities. Live migration can be performed to different physical systems, even across SPARC T-Series server generations. Encryption is employed automatically, with on-chip cryptographic acceleration, to ensure the security of the environment and data during the migration process. This is similar to AIX LPAR live migration. While Oracle Solaris Zones cannot be migrated while active, they can be cloned or easily moved from one global zone to another (even on other hardware systems).

#### Key Features

Oracle Solaris Cluster extends Oracle Solaris to provide enhanced availability of hosted applications. Using the advanced capabilities in Oracle Solaris, Oracle Solaris Cluster offers:

- **Flexible configurations.** Oracle Solaris Cluster supports pair, pair+N, N*1, N*N for flexible topologies, as well as clustering support for Oracle Solaris Zones.
- **Global devices, files, and networking.** All global devices, files, and network interfaces can be seen as local resources. Cluster nodes can access and utilize devices that are attached to another node within the cluster. These facilities create improved resource availability and simplified administration.
- **Virtualization support.** Oracle Solaris Cluster supports Oracle's virtualization portfolio—Oracle Solaris Zones, Oracle VM Server for SPARC (available on SPARC T-Series servers), and Dynamic Domains (available on SPARC Enterprise M-Series servers)—for flexible configurations that support consolidation efforts. Applications can run unmodified in virtualized environments.
- **Flexible storage support.** Oracle Solaris Cluster deployments can take advantage of a wide range of storage technologies, such as Fibre Channel, SCSI, iSCSI, and NAS storage solutions from Oracle and other vendors. Support for a broad range of file systems eases the data migration process.
- **Oracle RAC 10*g* and 11*g* integration and administration.** Automated installation and wizard-led configuration enable faster setup of Oracle RAC with Oracle Solaris Cluster. Specific Oracle RAC integration points enable improved coordination and simplified administration.
- **Campus and geographic clusters.** Oracle Solaris Cluster supports the creation of clusters across a campus or metropolitan area (campus cluster) or over large distances (geographic cluster) to support multi-site disaster recovery.

For more information on Oracle Solaris Cluster, see http://www.oracle.com/technetwork/server-

storage/solaris-cluster/overview/index.html and the Oracle Solaris Cluster Product Documentation at http://www.oracle.com/technetwork/server-storage/solaris-cluster/documentation/index.html.

## Differences Between AIX Clustering and Oracle Solaris Cluster

Table 5-2 summarizes the key differences between the combination of IBM clustering products and Oracle Solaris Cluster.

**TABLE 5-2. COMPARISON SUMMARY OF AIX CLUSTERING FACILITIES AND ORACLE SOLARIS CLUSTER**

| ITEM | AIX | ORACLE SOLARIS CLUSTER |
|---|---|---|
| Configuration | • 2 to 16 nodes<br>• Active/active, active/standby | • 2 to 16 nodes (SPARC), 2 to 8 (x86)<br>• Active/active, active/standby, rolling standby<br>• Pair, pair+N, N*1, N*N<br>• Oracle Solaris Zones |
| Networking Protocols | • IPv4 (1 required per node), IPv6 | • IPMP, Trunking, Jumbo Frames, VLAN<br>• IPv4, IPv6, SCTP, RDS |
| Disk Fencing | • Yes | • Yes |
| File Systems | • JFS2 and GPFS | • Root: Oracle Solaris ZFS<br>• Failover: UFS, Oracle Solaris ZFS, NFS<br>• Cluster: PxFS, Oracle ASM<br>• Support for Cluster File System (ACFS) and QFS |
| Volume Management | • AIX Logical Volume Manager | • Oracle Solaris ZFS<br>• Oracle Automatic Storage Management (ASM)<br>• Oracle Solaris Volume Manager |
| Virtualization Support | • LPARs<br>• WPARs<br>• Micro-Partitions | • Oracle Solaris Zones<br>• Oracle VM Server<br>• Dynamic Domains (on supported systems) |
| Monitoring | • System (heartbeat)<br>• Network<br>• Application<br>• PowerVM HMC | • System (heartbeat)<br>• Network<br>• Application<br>• Quorum<br>• Disk path<br>• Storage resources<br>• Oracle Solaris Cluster Web based GUI |
| Workload Management | • Yes | • Yes |

| | | |
|---|---|---|
| Cluster Management | • PowerHA<br>• CAA (AIX built in) | • Configuration Wizards<br>• Object-oriented command line interface<br>• Integrated with Oracle Enterprise Manager Ops Center<br>• Integrated with SMF |

# Management and Monitoring

Proactive monitoring and management help to prevent system errors and faults from impacting application response time and user productivity. By closely monitoring performance metrics and system health, administrators can take action before a problem escalates, service delivery deteriorates, and unplanned downtime occurs. AIX administrators rely heavily on SMIT and AHAFS. In heterogeneous environments, some AIX administrators use the Tivioli tool suite, which has long been supported on Oracle Solaris.

### Infrastructure Management in Oracle Solaris

For the Oracle Solaris administrator, Oracle offers an integrated set of tools (Oracle Enterprise Manager product line[13]) to provide end-to-end IT management that extends from applications to systems, virtual machines, software, and storage for both Oracle Solaris and Linux. The family of Oracle Enterprise Manager products enables management of the entire Oracle stack. Rich monitoring features support proactive application and systems management across the infrastructure for both Oracle and non-Oracle components.

A key part of the product family is Oracle Enterprise Manager Ops Center, which controls data center assets and simplifies physical and virtual server lifecycle management. The Ops Center software enables provisioning, patching, monitoring, administration, and configuration management via a Web-based user interface. As a result, it helps to reduce the complexity and cost associated with managing Oracle Solaris, Oracle VM Server, Linux, UNIX, and Windows operating system environments. The tool helps administrators to gain insight into Oracle server, storage, and network components, helping them to manage large numbers of systems in a more scalable manner.

Oracle monitoring is handled entirely via software interfaces. There is no direct equivalent to the maze of IBM hardware monitors, such as HMC, CMM, FSM, and IVM, or the short lived, software-based SDMC.

### Oracle Enterprise Manager Product Family

The Oracle Enterprise Manager family of products provides comprehensive solutions for testing, deploying, operating, monitoring, diagnosing, and resolving problems in complex IT environments. Administrators can manage the entire application life cycle with comprehensive application quality management and compliance solutions.

- **Cloud Lifecycle Management.** Oracle Enterprise Manager is a complete cloud management solution that includes self-service provisioning using centralized, policy-based resource management, integrated

[13] Oracle Enterprise Manager Ops Center licensing is included with Oracle Premier Support.

chargeback and capacity planning, and visibility into the physical and virtual environment from applications to disk.

- **Application Management.** Oracle Enterprise Manager provides the most complete management solution for Oracle Fusion Middleware applications, Oracle E-Business Suite, and Oracle's Siebel, PeopleSoft, and JD Edwards applications. It provides unique capabilities such as real user monitoring, zero-overhead instrumentation, and testing accelerators.
- **Middleware Management.** Delivering capabilities such as production diagnostics, model-driven topology mapping and business transaction management, Oracle Enterprise Manager is an end-to-end middleware management solution for Oracle Fusion Middleware environments.
- **Database Management.** From a single console, Oracle Enterprise Manager enables database manageability with real-time Automatic Database Diagnostic Monitor (ADDM) and Active Session History (ASH) analytics.
- **Application Performance Management.** Oracle Enterprise Manager provides a complete Application Performance Management (APM) solution for custom and Oracle applications, including Oracle E-Business Suite, Oracle Fusion Middleware, and Oracle's Siebel, PeopleSoft, and JD Edwards applications.
- **Application Quality Management.** Oracle Application Quality Management products provide a complete testing solution for Oracle Database, Oracle packaged applications, and custom Web applications.
- **Lifecycle Management.** Lifecycle Management is a comprehensive solution that helps database, system, and application administrators automate the processes needed to manage the lifecycle of Oracle technologies. It eliminates the manual and time-consuming tasks related to discovery, initial provisioning, patching, configuration management, and ongoing change management. In addition, it provides compliance frameworks for reporting and managing industry and regulatory compliance standards.
- **Engineered Systems Management.** Oracle Enterprise Manager supports comprehensive and centralized health monitoring and management of hardware and software components in Oracle's engineered systems, including Oracle SPARC SuperCluster T4-4, Oracle Exadata Database Machine, and Oracle Exalogic Elastic Cloud.
- **Hardware and Virtualization Management.** Described in more detail below, Oracle Enterprise Manager Ops Center combines the management of servers, operating systems, firmware, virtual machines, storage, and network fabrics into a single console. It delivers an integrated solution for physical and virtual server lifecycle management, providing comprehensive provisioning, patching, monitoring, administration, and configuration management.
- **Heterogeneous Management.** Oracle Enterprise Manager provides an extensible and customizable IT management framework. System Monitoring Plug-ins, developed both by Oracle and third-party vendors, add visibility into the underlying database, middleware, applications software, and hardware components. Oracle Enterprise Manager can be integrated with legacy third-party management tools and help desk systems, either by forwarding monitoring events or through the customized actions of Management Connectors. Management Connectors provide the bi-directional exchange of alerts, automatic help desk

ticket creation, and seamless workflow for incident management and resolution with other management tools. Oracle Enterprise Manager also supports AIX.[14]

### Oracle Enterprise Manager Ops Center

Designed to simplify the management of large numbers of IT systems, Oracle Enterprise Manager Ops Center is a comprehensive tool for managing Oracle Solaris, Oracle VM Server, Linux, and Windows servers—from firmware, server status, energy use, and operating systems, to virtual machines, storage, and network fabrics. Administrators use a centralized dashboard to gain cross-stack visibility into infrastructure resources (Figure 5-2). Capabilities include asset discovery, firmware and operating system provisioning, automated patch and configuration management, virtualization management, and compliance reporting.

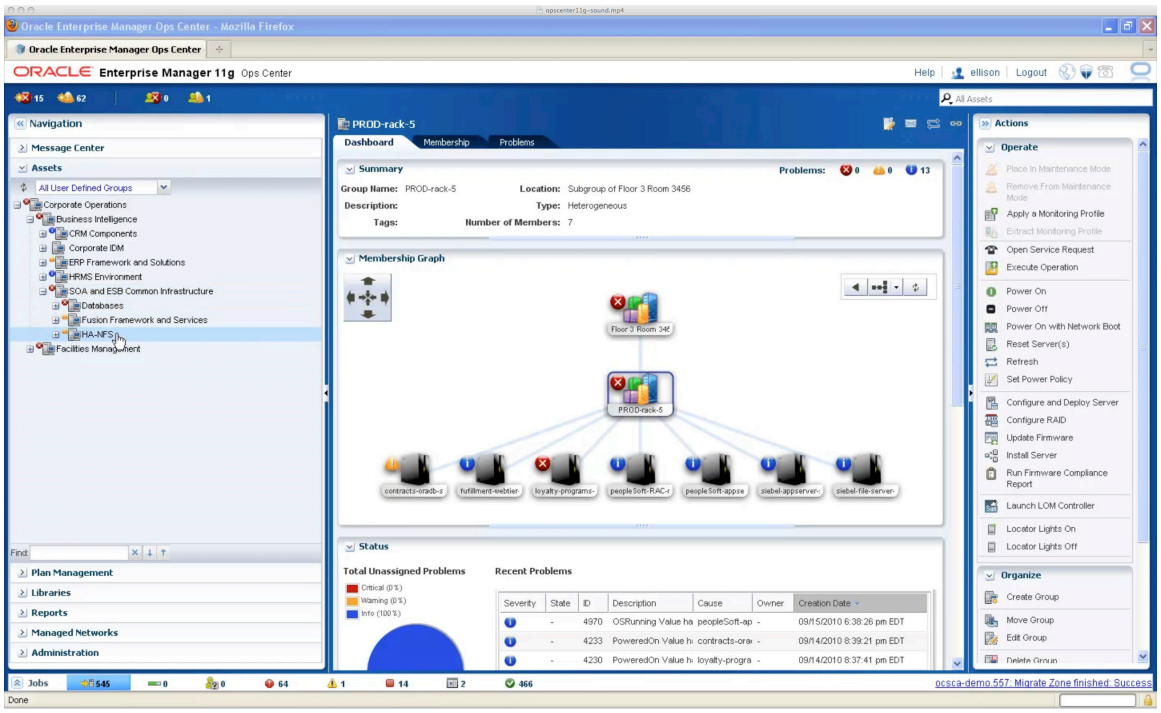


Figure 5-2. The Oracle Enterprise Manager Ops Center dashboard provides a comprehensive view of managed assets.

Oracle Enterprise Manager Ops Center includes these capabilities:

- **Systems monitoring.** Oracle Enterprise Manager Ops Center provides data about potential faults, performance metrics, and system health so that operational staff can resolve issues and take action, helping to prevent downtime.

[14] See http://www.oracle.com/technetwork/oem/grid-control/downloads/aix5lsoft-091214.html

- **Energy monitoring.** Modern servers use built-in service processors to monitor power consumption, making it possible for Oracle Enterprise Manager Ops Center to monitor and report on energy use for individual or logical groups of systems.
- **Operating system software monitoring.** Oracle Enterprise Manager Ops Center automatically installs lightweight software agents that track operating system performance characteristics, displaying collected data in easy-to-understand graphs and tables. Administrators can set site-specific thresholds to limit the amount of reported data, reducing administrative overhead and improving an operator's ability to detect problems.
- **Point-in-time provisioning.** Oracle Enterprise Manager Ops Center can discover new data center assets and access service processors on these systems to capture location identifiers, power states, network boot configurations, and firmware versions. Administrators can trigger firmware updates as needed. To expedite the deployment of new systems in the data center, Oracle Enterprise Manager Ops Center streamlines network configuration, automatically identifying MAC addresses, tracking ports, and maintaining this information in the DHCP configuration. Once the network is configured, Oracle Enterprise Manager Ops Center can use operating provisioning technologies such as the Oracle Solaris 11 Automated Installer, Oracle Solaris 10 JumpStart, Kickstart, or AutoYaST tools. Administrators can save common installation profiles and install new systems from them.
- **Software lifecycle management.** Operating system vendors commonly offer software updates to address security concerns, fix defects, or add new features. Oracle Enterprise Manager Ops Center creates dependency trees across software components, including trees for Oracle Solaris, Oracle Linux, Red Hat Enterprise Linux, Attachmate/Novell SUSE Linux, and Microsoft Windows distributions. By scanning patch readme files, package dependency files, RPM package manager header files, and other sources, Oracle Enterprise Manager Ops Center performs periodic baseline analysis on software update recommendations. Software updates to Oracle Solaris 11 take advantage of its Image Packaging System technology and Boot Environment capabilities, allowing administrators to update a clone of the active boot image, minimizing downtime for system upgrades and kernel package updates.
- **Infrastructure management.** Oracle Enterprise Manager Ops Center is a converged hardware management solution combining management across servers, operating systems, firmware, virtual machines, storage and network fabrics into a single console. Premier Support customers can use Oracle Enterprise Manager Ops Center at no additional charge.
- **Virtualization management.** Oracle Enterprise Manager Ops Center allows administrators to manage the full lifecycle of virtual machines defined with Oracle Solaris Zones, Oracle VM Server for x86, Oracle VM Server for SPARC, and even Dynamic Domains on the SPARC M-Series servers. Administrators can create, delete, stop, start, clone, and change the configuration of a virtual machine, and perform all of these actions on logical system groupings for efficient virtual machine management and provisioning.

For more information, see the Oracle Enterprise Manager Ops Center resources page at http://www.oracle.com/technetwork/oem/ops-center/index.html?ssSourceSiteId=ocomen.

## Infrastructure Monitoring

Oracle Solaris 11 provides a variety of monitoring tools that span different facets of the operating system.

- **Probes and tracing.** In AIX, the `probevue` dynamic trace facility provides system and application probes to enable administrators to observe live systems. Oracle Solaris includes the Oracle Solaris Dynamic Tracing (DTrace) facility, a dynamic tracing framework that provides top-to-bottom system observability for troubleshooting systemic problems in real time. Designed to quickly identify the root cause of system performance problems, DTrace combines over 100,000 trace points with a powerful scripting language and a simple, interactive command-line interface. It works by safely and dynamically instrumenting the running operating system kernel and applications with trace points (known as *probes*) that are completely passive until enabled. Probes can be enabled quickly for data collection, and then disabled again to minimize performance impacts on the system being examined. Developers and administrators can use this information to quickly identify performance bottlenecks, optimize resource utilization and performance, and quantify resource requirements.
- **Profiling.** Unifying application and system profiling on Oracle Solaris platforms, Oracle Solaris Studio DLight analyzes data from multiple sources in a synchronized fashion to trace and pinpoint application runtime problems. Incorporating the power of DTrace technology, Oracle Solaris Studio DLight enables developers to explore the system, understand how it works, and track down performance problems across many software layers. More importantly, remote capabilities make it possible for users to work at one system while monitoring services on another server running Oracle Solaris. An easy-to-use graphical interface provides application information, including thread microstates and data on CPU, memory, thread, and I/O usage for the duration of program execution.
- **Special-purpose tools.** DTrace provides a framework for building special-purpose performance monitoring tools. Using the toolkit included in the Oracle Solaris 11 repository, as well as developer tools such as the Performance Analyzer and DLight, developers can create tools that exploit the fundamental facilities provided by DTrace in the operating system kernel. For example, Oracle Solaris Studio DLight supports the creation of additional instruments, or *dtracelets*, for further observation of applications or systems. Dtracelets are XML files that are used to collect information and display specific data in particular ways. New dtracelets can be saved and displayed in the Oracle Solaris Studio DLight list of instruments and made available system-wide or used locally by individual developers.
- **Performance analysis.** In Oracle Solaris, the Performance Analyzer contains tools to help assess the performance of application code, identify potential performance problems, and locate the part of the code where the problems occur. Support is provided for MPI applications, including an MPI timeline and MPI charts, as well as zooming and filtering capabilities. Using Performance Analyzer in conjunction with a data Collector tool, developers can determine resource consumption levels, identify the functions that consume the most resources, find the source code lines and instructions that are responsible, and more.

### Using DTrace in Virtualized Environments

DTrace can be used in a zone to examine applications, identify performance bottlenecks, and quantify application resource requirements. By running DTrace in the global zone, system administrators can obtain a global (inter-zonal) view of what is happening in a production system. This allows for debugging subtle resource misallocation issues, a task that is much more difficult when each virtual system runs in isolation.

Using the `zonestat` command, system administrators can monitor CPU, memory, and network utilization, compare utilization rates to resource control limits, and determine how these resources are used in the zone over specified periods of time.

**TABLE 5-3. COMPARISON OF MONITORING TOOLS**

| TASK | IBM AIX | ORACLE SOLARIS |
|---|---|---|
| Monitoring (Primary tool) | • `probevue` command<br>• Dynamic tracing facility<br>• Uses the vue scripting language to determine where, what, and when to trace<br>• Trace points inserted at runtime | • DTrace<br>• Many providers, 100,000+ probes within the operating system stack<br>• Providers cover different aspects of the system<br>• Ability to observe individual processes<br>• Support for a variety of networking protocols<br>• Support for Java, Python, PHP, and Ruby runtimes |
| Monitoring (Other tools) | • hpmstat, iostat, netstat, procmon, sar<br>• svmon, top, topas, vmstat, xmstat | • Flowstat, dlstat, netstat, actadm, zonestat, svcs, fmstat<br>• iostat, fsstat, stat, kstate, mpstat, pgstat, vmstat, prstat<br>• truss, ptree, poolstat, latencytop, powertop |
| Hardware Management | • IBM Systems Director<br>• Basic hardware management through advanced cross-enterprise management (optional) | • Oracle Enterprise Manager Ops Center<br>• Included with all Oracle Premier Support agreements<br>• Extensive monitoring at enterprise scale |

# Chapter 6. Securing Infrastructure

AIX and Oracle Solaris both provide a comprehensive set of security features for enterprise-grade UNIX environments. Both environments provide role-based access controls (RBAC) that provide fine-grained control of permissions as well as trusted extensions for labeling and containing data when working in highly sensitive and secure environments. Oracle Solaris takes a defense-in-depth approach by providing tools to protect hosts, networks and data at rest. The full suite of security controls combined with the extensive virtualization capabilities in Oracle Solaris give IT organizations the tools needed to safely isolate applications and data in multi-tenant environments such as private and public clouds. Table 6-1 maps the security capabilities of AIX 7.1 to counterpart security features in Oracle Solaris.

**TABLE 6-1. SECURITY MAPPINGS**

| FUNCTIONALITY | AIX | ORACLE SOLARIS |
|---|---|---|
| Fine-grained privilege management | • Role-based access control | • Role-based access control<br>• Root account removed, root privileges are role-based |
| System hardening | • AIX Secure by default installation option<br>• AIX security expert, file system permission tool<br>• Trusted Execution | • Oracle Solaris secure by default configurations<br>• Basic Audit Reporting Tool (BART)<br>• Cryptographically signed package management |
| On-disk encryption | • JFS2 encryption | • Oracle Solaris ZFS dataset encryption<br>• Oracle Solaris Cryptographic Framework |
| Trusted Computing | • Trusted AIX | • Oracle Solaris Trusted Extensions<br>• Labeled IPsec, labeled Oracle Solaris ZFS datasets |
| Firewall | • AIX TCP/IP Filter | • IP Filter with SMF integration |
| Application isolation and delegated administration | • Workload Partitions (WPARs)<br>• Domain-based role-based access control | • Oracle Solaris Zones, delegated administration<br>• Optional private IP stack/zone (no dedicated hardware) |
| Secure remote login | • Secure Shell (`ssh`) | • Secure Shell with X.509 certificate extensions |
| Common Criteria requirements level[15] | • Common Criteria at Common Access Protection Profile (not yet certified)<br>• CAPP/Evaluation Assurance Level (EAL) 4+, including the Role-Based Access Control Protection Profile (RBACPP), and the Labeled Security Protection Profile (LSPP) | • Common Criteria at Common Access Protection Profile<br>• CAPP)/Evaluation Assurance Level (EAL) 4+, including the Role-Based Access Control Protection Profile (RBACPP), and the Labeled Security Protection Profile (LSPP) |

[15] In evaluation status. See http://www.oracle.com/technetwork/topics/security/security-evaluations-099357.html

More information on these security topics can be found in the *Oracle Solaris 11 Security Guidelines* manual at http://docs.oracle.com/cd/E23824_01/html/819-3195/index.html and the Oracle Solaris Administration: Security Services manual at http://docs.oracle.com/cd/E23824_01/html/821-1456/index.html.

## Role-Based Access Control

In the past, UNIX systems provided only the most basic access controls. The traditional superuser (`root`) account enabled completely unrestricted access to the whole system. Anyone with the root password could log in directly, leaving no audit trail or indication of who is using this wide-open access mechanism. Often, even the most junior administrators were given the root access, because there was no fine-grained mechanism for providing only the access that was required for the job.

Both AIX 7.1 and Oracle Solaris 11 offer powerful user and process rights management technology that reduces security risks by granting users and applications only the minimum capabilities needed to perform tasks. Process rights management enables processes to be restricted at the command, user, role, or system level. In Oracle Solaris, process rights management is implemented through a privilege mechanism. Granting only the least necessary privileges reduces the security risk compared to a user or process having full superuser capabilities on a system.

Privileges and role-based access control (RBAC) provide a compelling alternative to the traditional superuser model. In the RBAC model on Oracle Solaris, users log in as themselves and assume roles when needed that enable them to perform tasks that require elevated privileges. In Oracle Solaris 11, even root access follows this model. By making root a role, users cannot log in as root directly. Administrators must first log in under their regular user ID. When root privileges are needed, administrators must enter a password to assume the root role. Unlike becoming the superuser on traditional systems, the identity of the user remains the same. (It does not change to root.) Assuming the root role temporarily adds the privileges associated with the root role to the person's existing account. Removing direct access to the root account provides better audit trails that are tied to a specific user and makes the system more secure. While AIX 7.1 offers enhanced RBAC that gives the option to define privileged roles, Oracle Solaris 11 has taken the additional step to remove the root account itself.

Most UNIX systems include a setuid mechanism that can be set on a command to enable any user who runs it to run as specific user (usually root) to perform a specific task. A setuid root executable has complete access to the system and network when the command is running. As a result, compromising a setuid executable can give an attacker unrestricted access to the system. Oracle Solaris uses a privilege framework that allows fine-grained control of what a process can do, such as read or write files, gain network access, or manipulate processes. These controls can be used to limit the capabilities granted to any privileged executable to only the specific rights that are needed for the task. This can significantly limit potential vulnerabilities associated with setuid executables.

More information on RBAC in Oracle Solaris can be found in the "Roles, Rights Profiles, and Privileges" section of the *Oracle Solaris Administration: Security Services* guide at http://docs.oracle.com/cd/E23824_01/html/821-1456/prbactm-1.html#scrolltoc.

## Host Security

AIX and Oracle Solaris include an extensive set of mechanisms for ensuring host security. Both systems offer a secure-by-default installation choice, improving security by installing reduced configurations and ensuring that a minimal set of network accessible services are available. Both systems also offer a number of tools for ensuring the integrity of the software installed on the system. These tools can detect whether system files have been accidentally or maliciously modified. Oracle Solaris 11 features include:

- **Basic Auditing Reporting Tool.** This tool validates systems by performing file-level checks of a system over time. It is similar to the file system permissions tool in AIX. By creating a baseline manifest for a newly installed and configured system, it is possible to gather information about installed software. This baseline information can be compared to a snapshot of the system at a later time. A generated report lists file-level changes that have taken place since system installation, enabling administrators to verify that no accidental or malicious changes have been made to the environment.
- **Image Packaging System**. The manifest maintained by the Imaging Packaging System describes all package files and their permissions. This information can be used to verify the integrity of installed packages by using the `pkg verify` command. In addition, Oracle Solaris 11 allows package authors to digitally sign packages, which can be used to ensure authenticity. (Administrators can make this an enforced policy, enabling only signed packages with valid signatures to be installed.) See http://docs.oracle.com/cd/E19963-01/html/820-6572/gkkos.html for more information.
- **Audit access**. The Oracle Solaris audit feature provides the ability to log system activity for any auditable Oracle Solaris event—such as system calls on the server machine, packets sent over the network, or a sequence of bits written to disk—at a granular level. Starting with Oracle Solaris 11, auditing is a service that is managed by SMF, and auditing records are stored in binary files on an Oracle Solaris ZFS file system. System reboots are not required to enable auditing.

See "Using the Basic Audit Reporting Tool" at http://docs.oracle.com/cd/E23824_01/html/821-1456/bart-1.html#scrolltoc, "Verifying Package Installation" at http://docs.oracle.com/cd/E23824_01/html/E21802/gkoks.html#gilmv, and "Monitoring SMF Services" at http://docs.oracle.com/cd/E23824_01/html/821-1451/dzhaq.html#scrolltoc for more information.

## Network Security

A key component of the defense-in-depth approach to security in Oracle Solaris is extensive network security. Designed to protect physical and virtual systems and networks, Oracle Solaris creates a fully virtualized and secure network that provides protection even if intruders gain access to physical or virtual systems.

- **Secure by default configuration.** When Oracle Solaris is installed, a large set of network services are disabled by default. In this secure-by-default configuration, the only network service that accepts network requests is the `sshd` daemon. All other network services are disabled or only handle local requests. These services are managed just like all other services, using SMF, simplifying administration. Individual network services, such as `ftp`[16], can be enabled through the SMF interface.
- **Integrated firewall.** Both AIX 7.1 and Oracle Solaris 11 include an integrated stateful packet filtering firewall. The Oracle Solaris firewall is based on the open source IP Filter, which provides basic firewall services and network address translation (NAT). In Oracle Solaris 11, the IP Filter software is configured and managed using SMF. The `pfil` module is replaced by packet filter hooks, streamlining the procedure to enable the IP Filter software. Through these hooks, IP Filter uses pre-routing (input) and post-routing (output) filter taps to control packet flow into and out of the Oracle Solaris system. The results are improved performance and the ability to filter traffic between Oracle Solaris Zones.
- **Secure Shell with X.509 certificate extensions.** Both AIX 7.1 and Oracle Solaris 11 include the Secure Shell (`ssh`) for secure remote login. In Oracle Solaris 11, the `ssh` utility is extended to allow the use of X.590 certificates for authenticating users and hosts. This makes the use and administration of the `ssh` utility more straightforward and secure. Users do not need to populate authorized key files on each host or answer prompts to verify a host's authenticity since it can be verified using the host's X.509 certificate.

For more information on the IP Filter implementation in Oracle Solaris, see the "IP Filter in Oracle Solaris (Overview)" section of the *Oracle Solaris Administration: IP Services* guide at http://docs.oracle.com/cd/E23824_01/html/821-1453/eupsq.html

## On-Disk Encryption

Keeping data secure when it is at rest is just as important as keeping it safe as it travels over networks. Using encryption to protect files can help protect information from being stolen or manipulated in the event of a network security breach. Just as AIX JFS2 provides an encrypted file system (EFS), Oracle Solaris ZFS supports encrypted datasets. The ability to encrypt datasets helps protect against the theft of devices, man-in-the-middle attacks on the SAN, and dataset-level secured deletion. In addition, Oracle Solaris ZFS is integrated with Oracle Solaris Cluster to support highly available, secure infrastructure deployments.

JFS2 and EFS support per-file encryption. In Oracle Solaris ZFS, data is encrypted at the dataset level rather than on a per-disk basis. This allows users to place a mix of encrypted and unencrypted information in the same storage pool. Whether or not to encrypt data in a dataset is determined at dataset creation. User data and

[16] This is not to suggest that enabling standard `ftp` is recommended. The `sftp` utility provides equivalent functionality, but with the same security as `ssh`. If traditional `ftp` is to be enabled, it is probably best to enable it in a tightly locked down immutable zone.

file system metadata are encrypted using a sophisticated encryption key management facility to support different key management strategies.

Oracle Solaris ZFS uses the cryptographic framework built into the operating system to enable cryptographic protection of data on a per-dataset basis. Based on the PKCS#11 public key cryptography standard created by RSA Security, Inc., the framework provides a mechanism and API whereby both kernel- and user-based cryptographic functions can transparently use software encryption modules and hardware accelerators configured on the system. Applications that run on Oracle's SPARC T-Series servers and use the cryptographic framework gain an added benefit: cryptographic operations are offloaded automatically to an on-chip cryptographic accelerator in the server for improved performance.

The framework provides various services, including message encryption and message digest, message authentication, and digital signing. It also includes APIs for accessing cryptographic services, and SPIs for providing cryptographic services. New cryptographic enhancements in Oracle Solaris 11 include support for FIPS 140-2 of the Federal Information Processing Standard, and the implementation of ECC and other NSA Suite B protocols to meet stringent government standards. These routines are highly optimized and are used automatically by Oracle Solaris ZFS, the Java Cryptography Extension (JCE), Kerberos, IPsec, and other components in Oracle Solaris.

More information on the cryptographic framework can be found in the "Cryptographic Services" section of the *Oracle Solaris Administration: Security Services* guide at http://docs.oracle.com/cd/E23824_01/html/821-1456/scftm-1.html#scrolltoc.

## Server Virtualization Security

The capability to delegate administration for individual zones without giving away administrative access to the host system is provided with Oracle Solaris delegated administration. This feature is similar to the domain support for RBAC introduced in AIX 7.1. These security features take on extra importance as IT organizations utilize cloud-based infrastructure, software, and services. Oracle Solaris builds virtualization and security into the operating system, making it an ideal platform for cloud-based infrastructure.

- **Delegated administration.** Similar to the domain support for RBAC in AIX 7.1, Oracle Solaris enables the delegation of common administrative tasks for specific zones to specific administrators using role-based access controls. This is particularly powerful in a shared environment where it is desirable to allow specific users to only manage zones that are relevant to their role.
- **Zone link protection.** In many virtualized environments, it is common for the host administrator to grant exclusive access of a physical link or a virtual network interface card (NIC) to a guest virtual machine. Doing so enables guests to benefit from traffic isolation and improved performance. On the other hand, guests can generate any type of traffic, even harmful packets, and send it over the network. Oracle Solaris 11 provides a new link protection mechanism for preventing potentially malicious or misbehaving guest virtual machines from sending harmful packets to the network. This feature provides protection against basic threats, including IP, DHCP, MAC, and L2 frame spoofing.

- **Exclusive IP zones.** In many virtualized environments, a dedicated physical network interface controller is needed to achieve network separation between virtual environments. Exclusive IP zones in Oracle Solaris give administrators the ability to assign a separate IP stack for each zone. The stack is completely separate from all other zones yet does not require the expense or complexity of a dedicated network connection.
- **Immutable zones (read-only root).** Oracle Solaris supports the creation of immutable zones (zones that have read-only file systems). Using mandatory write access control, read-only file systems cannot be modified by processes running in the zone, even those with root privileges. Writes can take place only in the system's global zone. This provides truly robust protection for application stacks running within immutable zones.

## Trusted Computing

Both AIX 7 and Oracle Solaris 11 include Trusted Extensions, an optionally installable layer of secure labeling technology that enables data security policies to be separated from data ownership. Using Trusted Extensions, access to data is controlled by special security tags. These tags are called labels. Labels are assigned to users, processes, and objects such as data files and directories. These labels are used to enforce mandatory access control (MAC) in addition to UNIX permissions, or discretionary access control (DAC).

The following highlight the key aspects of security extensions and how they work in Oracle Solaris.

- **Enforced policies and labels.** Both AIX and Oracle Solaris provide these facilities.
- **Labels and access control.** Trusted Extensions supports traditional discretionary access control (DAC) policies based on ownership, as well as label-based mandatory access control (MAC) policies. When the label-based MAC policies are enabled, all data flows are restricted based on a comparison of the labels associated with the processes (subjects) requesting access and the objects containing the data. Unlike most other multi-level operating systems, Oracle Solaris includes a multi-level desktop. (In AIX 7.x, the X environment is disabled when Trusted AIX is installed.)
- **Credentials.** To enable greater flexibility and security, the Trusted Extensions feature in Oracle Solaris 11 enables per-label and per-user credentials. As a result, administrators can require a unique password for each label. Since this password is in addition to the session login password, administrators can set a per-security zone encryption key for each label of every user's home directory.
- **Network communications.** When labeled processes in a multi-level secure operating system communicate across system boundaries, their network traffic must be labeled and protected. Typically this requirement is met by using physically separate network infrastructure to ensure that data belonging to different labeled domains stays in separate physical infrastructures. In Oracle Solaris 11, security labels on packets received from remote hosts and labeled IPsec/IKE enable organizations to reuse the same physical network infrastructure for labeled communications. This is accomplished by transferring labeled data within separate labeled IPsec security associations, removing the need for redundant and expensive physical network infrastructure.
- **Datasets.** In Oracle Solaris 11, Trusted Extensions enables security labels to be set on Oracle Solaris ZFS datasets. When used, the security labels ensure that Oracle Solaris ZFS file systems used for a specific

security label cannot be mounted on a physical or virtual system with a different security label. This restriction helps to avoid the inadvertent upgrade or downgrade of the classification of data.

- **Common Criteria Requirements.** Both AIX 7.1 and Oracle Solaris Trusted Extensions are designed to meet the requirements of the Common Criteria Labeled Security Protection Profile (LSPP), the Role-Based Access Protection Profile (RBACPP), and the Controlled Access Protection Profile (CAPP). What makes the Oracle Solaris implementation unique is its ability to provide high assurance while maximizing compatibility and minimizing overhead.
- **Trusted Platform Module.** AIX includes a software implementation of the Trusted Platform Module (TPM) that stores digital certificates, keys, and passwords. Based on TrouSerS, the open source Trusted Computing Group (TCG) Software Stack, TPM provides access to RSA key pair generation, encryption, decryption, and storage. Oracle Solaris provides similar functionality and utilizes hardware TPM modules when available. It includes the TrouSerS package, a PKCS#11 provider for using the TPM to store keys, and a `tpmadm` utility for performing administrative functions and viewing the state of TPM registers.

For more information, see the *Oracle Solaris Trusted Extensions User's Guide* at http://docs.oracle.com/cd/E23824_01/html/821-1484/index.html.

## Scripting

System administrators write scripts to manage as much of their day-to-day tasks as practical. AIX and Oracle Solaris both provide the `ksh93` shell. (AIX provides the "t" version while Oracle Solaris provides the "u" version.) Most scripts will port with minimal effort related to script dialect variation.

Many AIX administrators also rely on Rexx, another powerful scripting language. While Oracle Solaris does not ship Rexx, open source and free for download versions are widely available and run well on Oracle Solaris. The open source tools include Regina (http://regina-rexx.sourceforge.net/), NexRexx (http://www-01.ibm.com/software/awdtools/netrexx/), and Rexx/imc (http://www.cs.ox.ac.uk/people/ian.collier/Rexx/rexximc.html). Regina is a native Rexx environment. Binaries and source are available for download for Oracle's x86 and SPARC platforms. NetRexx is a Java-based Rexx environment that was authored by IBM.

# Chapter 7. For More Information

Additional information and resources can be found in the references listed in Table 7-1.

**TABLE 7-1. ADDITIONAL READING**

| | |
|---|---|
| **ORACLE SOLARIS** | |
| Oracle Solaris | http://www.oracle.com/solaris |
| Oracle Solaris Technical Information | http://www.oracle.com/technetwork/server-storage/solaris11/ |
| Oracle Solaris 11 Technology Spotlights | http://www.oracle.com/technetwork/server-storage/solaris11/technologies/index.html |
| Oracle Solaris 11 Documentation | http://www.oracle.com/technetwork/server-storage/solaris11/documentation/index.html |
| Oracle Solaris 11 How-To Articles | http://www.oracle.com/technetwork/server-storage/solaris11/documentation/how-to-517481.html |
| *Oracle Solaris 11 What's New* | http://www.oracle.com/technetwork/server-storage/solaris11/documentation/solaris11-whatsnew-201111-392603.pdf |
| *Oracle Solaris 11.1—What's New* | http://www.oracle.com/technetwork/server-storage/solaris11/documentation/solaris11-1-whatsnew-1732377.pdf |
| **DATA MANAGEMENT** | |
| *Oracle Solaris Administration: Devices and File Systems* | http://download.oracle.com/docs/cd/E23824_01/html/821-1459/index.html |
| *Oracle Solaris Administration: ZFS File Systems* | http://download.oracle.com/docs/cd/E23824_01/html/821-1448/index.html |
| *Oracle Solaris Administration: SAN Configuration and Multipathing* | http://docs.oracle.com/cd/E23824_01/html/E23097/index.html |
| **HIGH AVAILABILITY AND SYSTEM MANAGEMENT** | |
| *Oracle Solaris 11 Administration: Common Tasks* | http://docs.oracle.com/cd/E23824_01/html/821-1451/toc.html |
| Oracle Solaris Cluster | http://www.oracle.com/us/products/servers-storage/solaris/cluster/overview/index.html |
| Oracle Solaris Cluster Technical Information | http://www.oracle.com/technetwork/server-storage/solaris-cluster/overview/index.html |
| Oracle Solaris Cluster Documentation | http://www.oracle.com/technetwork/server-storage/solaris-cluster/documentation/index.html |
| Oracle Solaris Cluster Training | http://www.oracle.com/technetwork/server-storage/solaris- |

| | |
|---|---|
| | cluster/training/index.html |
| Oracle Enterprise Manager 12*c* | http://www.oracle.com/technetwork/oem/grid-control/overview/index.html |
| Oracle Enterprise Manager Extensions Exchange | http://www.oracle.com/technetwork/oem/extensions/index.html |
| Oracle Enterprise Manager 12*c* Cloud Control Documentation | http://docs.oracle.com/cd/E24628_01/index.htm |
| **NETWORKING** | |
| *Oracle Solaris Administration: IP Services* | http://docs.oracle.com/cd/E23824_01/html/821-1453/index.html |
| *Oracle Solaris Administration: Naming and Directory Services* | http://docs.oracle.com/cd/E23824_01/html/821-1455/index.html |
| *Oracle Solaris Administration: Network Interfaces and Network Virtualization* | http://docs.oracle.com/cd/E23824_01/html/821-1458/index.html |
| *Oracle Solaris Administration: Network Services* | http://docs.oracle.com/cd/E23824_01/html/821-1454/index.html |
| **SECURITY** | |
| *Oracle Solaris 11 Security Guidelines* | http://docs.oracle.com/cd/E23824_01/html/819-3195/index.html |
| *Oracle Solaris Administration: Security Services* | http://docs.oracle.com/cd/E23824_01/html/821-1456/index.html |
| *Oracle Solaris Trusted Extensions User's Guide* | http://docs.oracle.com/cd/E23824_01/html/821-1484/index.html |
| *Trusted Extensions Configuration and Administration* | http://docs.oracle.com/cd/E23824_01/html/821-1482/index.html |
| *Trusted Extensions Label Administration* | http://docs.oracle.com/cd/E23824_01/html/821-1481/index.html |
| **SERVICES** | |
| "How to Create an Oracle Solaris Service Management Facility Manifest" white paper | http://www.oracle.com/technetwork/server-storage/solaris/solaris-smf-manifest-wp-167902.pdf |
| "Management of Systems and Services Made Simple with the Oracle Solaris Service Management Facility" white paper | http://www.oracle.com/technetwork/server-storage/solaris/solaris-smf-wp-167901.pdf |
| "SMF Concepts" | http://docs.oracle.com/cd/E23824_01/html/821-1451/dzhid.html |
| **SOFTWARE MANAGEMENT** | |
| *Installing Oracle Solaris 11 Systems* | http://docs.oracle.com/cd/E23824_01/html/E21798/index.html |
| *Introduction to Managing Boot Environments* | http://docs.oracle.com/cd/E23824_01/html/E21801/index.html |
| *Creating a Custom Oracle Solaris 11 Installation* | http://docs.oracle.com/cd/E23824_01/html/E21800/index.html |

| | |
|---|---|
| *Image* | |
| *Adding and Updating Oracle Solaris 11 Software Packages* | http://docs.oracle.com/cd/E23824_01/html/E21802/index.html |
| *Copying and Creating Oracle Solaris 11 Package Repositories* | http://docs.oracle.com/cd/E23824_01/html/E21803/index.html |
| *Oracle Solaris 11 Installation Man Pages* | http://docs.oracle.com/cd/E23824_01/html/E21797/index.html |
| *Image Packaging System Man Pages* | http://docs.oracle.com/cd/E23824_01/html/E21796/index.html |
| **VIRTUALIZATION** | |
| *Oracle Solaris Administration: Oracle Solaris Zones, Oracle Solaris 10 Zones, and Resource Management* | http://docs.oracle.com/cd/E23824_01/html/821-1460/index.html |
| *Oracle Solaris Administration: Network Interfaces and Network Virtualization* | http://docs.oracle.com/cd/E23824_01/html/821-1458/index.html |
| **OTHER ADMINISTRATION** | |
| Evaluating Oracle Solaris 11 (All resources) | http://www.oracle.com/technetwork/server-storage/solaris11/overview/evaluate-1530234.html |
| *Oracle Solaris Tunable Parameters Reference Manual* | http://docs.oracle.com/cd/E23824_01/html/821-1450/index.html |
| *Oracle Solaris Administration: SMB and Windows Interoperability* | http://docs.oracle.com/cd/E23824_01/html/821-1449/index.html |
| "How to Perform System Archival and Recovery Procedures with Oracle Solaris 11" | http://www.oracle.com/technetwork/articles/servers-storage-admin/o11-091-sol-dis-recovery-489183.html |
| *International Language Environment Guide* | http://docs.oracle.com/cd/E23824_01/html/E26033/index.html |
| **TRAINING** | |
| Oracle Solaris 11 Training and Support | http://www.oracle.com/technetwork/server-storage/solaris11/training/index.html |
| Oracle University | http://education.oracle.com/ |

# Appendix A. Glossary

**Access Control List (ACL)**
A file containing a list of principals with certain access permissions. Typically, a server consults an access control list to verify that a client has permission to use its services.

**Agent Builder**
A component of Oracle Solaris Cluster that automates the creation of data services.

**API**
Application programming interface.

**appcert**
A utility that examines an application's conformance to the Oracle Solaris Application Binary Interface. Use of the `appcert` utility can help identify potential binary compatibility issues when porting applications to Oracle Solaris.

**Authentication**
A security service that verifies a claimed identity.

**Authorization**
The process of determining whether a user can use a service, which objects the user can access, and the type of access allowed.

**Big Endian**
An architecture that stores the most-significant byte of data first. Oracle Solaris uses a Big Endian architecture on SPARC processor-based systems and a Little Endian architecture on x86 platforms. AIX also uses a Big Endian architecture.

**Chip-Multithreading Technology**
Multithreaded processor technology that enables each processor core to switch between multiple threads on each clock cycle.

**CMT**
See Chip-Multithreading Technology.

**Data transformation**
The process of converting data from one format to another.

**DTrace**
See Oracle Solaris DTrace.

**ETL utilities**
Extract, Transform, and Load utilities, tools that take a wide array of formats and convert them into Structured Query Language (SQL) for relational database management systems.

**Hybrid Storage Pool**
A combination of disk drives and Flash devices that work together to minimize the impact of disk latencies and improve application performance. Flash devices handle certain types of I/O while hard disk drives store massive data sets. Hybrid Storage Pools are enabled by Oracle Solaris ZFS.

**Korn shell**
Oracle Solaris 11 provides `ksh93` (version u) as the default shell (located in `/usr/bin/ksh`), as well as `ksh88` (located in `/usr/sunos/bin/ksh`) for compatibility. AIX 7.1 provides the `ksh93t` shell.

**Little Endian**
An architecture that stores the least-significant byte of data first. Oracle Solaris uses a Little Endian architecture on x86 systems and a Big Endian architecture on SPARC processor-based platforms.

**Oracle Solaris Cluster**
A high availability solution for Oracle Solaris that is integrated at the kernel level. It monitors servers, storage, network components, operating system, virtual machines, and applications. Recovery actions are based on policies and application specifications.

**Oracle Solaris Cryptographic Framework**
A framework built into Oracle Solaris that provides kernel-level and user-level consumers with access to software-based or hardware-based cryptographic capabilities.

**Oracle Solaris DTrace**
A dynamic tracing facility built into Oracle Solaris that lets developers observe operating system and application behavior in real time.

**Oracle Solaris Key Management Framework**
A framework that provides tools and programming interfaces for managing PKI objects.

**Oracle Solaris Service Management Facility**
A facility introduced in Oracle Solaris 10 to simplify service management and control.

**Oracle Solaris Studio**
A free, comprehensive C, C++, and Fortran tool suite for Oracle Solaris and Linux operating systems that accelerates the development of scalable, secure, and reliable enterprise applications.

**Oracle Solaris ZFS**
A 128-bit file system that integrates volume management and provides virtually unlimited file system scalability.

**Oracle VM Server**
Scalable server virtualization software that supports Oracle and non-Oracle applications. Oracle VM Server is available for Oracle's SPARC and x86 servers.

**Package**
A collection of files and directories required for a software product. In Oracle Solaris, applications are distributed for deployment in packages.

**PKI**
Public Key Infrastructure.

**POSIX**
Portable Operating System Interface for UNIX. A set of standards that provide a well-defined system call interface for kernel facilities, as well as shell and utility interfaces.

**Privilege**
A discrete right that can be granted to an application.

**SMF**
See Oracle Solaris Service Management Facility.

**Trusted Extensions**
An optional layer of secure label technology in Oracle Solaris that allows data security policies to be separated from data ownership. Multilevel data access policies support compliance goals.

**UFS**
UNIX File System. In Oracle Solaris 11, UFS no longer is the default file system. Support for existing file UFS systems remains, and new UFS file systems can be created as needed.

**ZFS**
See Oracle Solaris ZFS.

IBM AIX to Oracle Solaris
Technology Mapping Guide
January 2013

Oracle Corporation
World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065
U.S.A.

Worldwide Inquiries:
Phone: +1.650.506.7000
Fax: +1.650.506.7200

oracle.com

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



**Hardware and Software, Engineered to Work Together**