

D82701GC10

Edition 1.0

December 2013

D85049

ORACLE®

Oracle Solaris 11 System Administration for Experienced UNIX/Linux Administrators

Student Guide • Volume II

Authors

Uma Sannasi
J S Raghavendra

**Technical Contributors
and Reviewers**

Geetha Nazare
Shripad Patki
John Hathaway
Harry Burks
Rosemary Martinak
Todd Lowry
Rajesh Rajasekharan
Gary Riseborough
David S Maxwell
Glynn Foster
Alex Barclay
Deepak Dhanukodi
Yousuf Mohammed
Niveditha Sananth
Sharvani Nagamalli
Pranamya Jain
Sreejith Mohan
Manish Pawar

Editors

Daniel Milne
Richard Wallis
Smita Kommini

Graphic Designers

Prakash Damodaran
Seema Bopaiah

Publishers

Syed Ali
Srividya Rameshkumar

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Disclaimer

This document contains proprietary information and is protected by copyright and other intellectual property laws. You may copy and print this document solely for your own use in an Oracle training course. The document may not be modified or altered in any way. Except where your use constitutes "fair use" under copyright law, you may not use, share, download, upload, copy, print, display, perform, reproduce, publish, license, post, transmit, or distribute this document in whole or in part without the express authorization of Oracle.

The information contained in this document is subject to change without notice. If you find any problems in the document, please report them in writing to: Oracle University, 500 Oracle Parkway, Redwood Shores, California 94065 USA. This document is not warranted to be error-free.

Restricted Rights Notice

If this documentation is delivered to the United States Government or anyone using the documentation on behalf of the United States Government, the following notice is applicable:

U.S. GOVERNMENT RIGHTS

The U.S. Government's rights to use, modify, reproduce, release, perform, display, or disclose these training materials are restricted by the terms of the applicable Oracle license agreement and/or the applicable U.S. Government contract.

Trademark Notice

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

Contents

Preface

1 Course Introduction

- Overview 1-2
- Course Goals 1-3
- Skills Gained 1-4
- Course Agenda: Day 1 1-5
- Course Agenda: Day 2 1-6
- Course Agenda: Day 3 1-7
- Course Agenda: Day 4 1-8
- Course Agenda: Day 5 1-9
- How Prepared Are You? 1-10
- Introductions 1-11
- Your Learning Center 1-12
- Your Lab Environment 1-13

2 Administering System Software by Using IPS

- Objectives 2-3
- Lesson Agenda 2-4
- Oracle Solaris: The Operating System 2-5
- Supported Platforms 2-6
- Software Management Prior to Oracle Solaris 11 2-7
- Software Management in Oracle Solaris 11 2-8
- Overview of IPS 2-9
- Publishers, Packages, and Repositories 2-10
- FMRI 2-11
- Repository Origins and Mirrors 2-12
- Images and BEs 2-13
- IPS Commands 2-14
- Quiz 2-15
- Lesson Agenda 2-16
- Why Do You Need a Local Repository? 2-17
- What Is a Local Repository? 2-18
- Local Repository Configuration Options 2-19
- Configuring a Local IPS Repository by Using SMF 2-20

Quiz	2-23
Lesson Agenda	2-24
Configuring Client Access to the Local IPS Server	2-25
Verifying Prerequisites Setup	2-26
Setting the Local IPS Publisher	2-27
Testing Client Access to the Local IPS Server	2-28
Lesson Agenda	2-29
IPS Interfaces	2-30
Package Management: CLI	2-31
Package Management: Package Manager	2-32
Lesson Agenda	2-33
Software Updates	2-34
Software Update Process	2-35
Update Interfaces	2-36
Updating Your Local Repository	2-37
Updating the OS by Using the CLI	2-38
Updating the OS by Using the GUI	2-40
Quiz	2-41
Lesson Agenda	2-42
Upgrading Oracle Solaris 11 to Oracle Solaris 11.1 OS	2-43
Upgrading the OS by Using the Oracle Solaris Support Repository	2-44
Upgrading the OS by Using the Oracle Solaris Release Repository	2-49
Quiz	2-51
Lesson Agenda	2-52
Overview of Boot Environments	2-53
Boot Environment Creation	2-54
BE Management Utilities	2-55
Administering Boot Environments	2-56
Listing the BEs on the System	2-57
Creating a New BE	2-58
Renaming an Existing Inactive BE	2-59
Destroying an Existing Inactive BE	2-60
Activating an Existing Inactive BE	2-61
Verifying the New BE	2-62
Mounting an Inactive BE	2-63
Unmounting an Inactive BE	2-64
Installing a Package on an Inactive Mounted BE	2-65
Uninstalling a Package on an Inactive Mounted BE	2-66
Creating a Backup of a BE	2-67
Creating a BE From an Existing Backup	2-68
Managing BEs with Package Manager	2-69

Quiz 2-70
Summary 2-71

3 Administering Services by Using SMF

Objectives 3-3
Lesson Agenda 3-4
Managing Services in Older UNIX OSs 3-5
Managing Services Since Oracle Solaris 10 3-6
Overview of SMF 3-7
SMF Concepts 3-8
SMF Service 3-9
Service Identifier 3-10
Service States 3-11
SMF Components 3-12
SMF Profile 3-13
When Are SMF Profiles Applied? 3-14
SMF Profile: Example 3-15
Service Configuration Repository 3-16
SMF Master Restarter Daemon (svc.startd) 3-17
Milestone 3-18
SMF Manifests 3-19
SMF Manifest: Example 3-20
SMF Repository Backups 3-21
SMF Repository Snapshots 3-22
Quiz 3-23
Lesson Agenda 3-24
Administering SMF Services 3-25
Listing Services Information 3-26
Displaying the Status of a Service Instance 3-27
Displaying the Service Dependents and Dependencies 3-28
Disabling a Service 3-29
Enabling a Service 3-30
Refreshing and Restarting a Service 3-31
Managing SMF Services Properties 3-32
Modifying inetd Service Properties 3-33
Managing SMF Services by Using the GUI 3-34
Quiz 3-35
Lesson Agenda 3-36
Configuring SMF Services 3-37
Creating a Service 3-38
Creating a Service: Example 3-39

- Creating a Service by Using svcbundle 3-43
- Modifying a Service's Manifest 3-44
- Changing an Environment Variable of a Service 3-45
- Changing a Property for an inetd-Controlled Service 3-46
- Creating and Applying an SMF Profile 3-47
- Changing Services and Their Configurations by Using the net services
Command 3-48
- Setting Up Service State Transition Notifications 3-49
- Managing Notifications 3-51
- Lesson Agenda 3-52
- Least Privilege and SMF 3-53
- Service Privileges 3-54
- SMF Rights Profile 3-55
- Authorizations and Rights 3-56
- Service-Specific Property Groups 3-57
- Quiz 3-58
- Lesson Agenda 3-59
- Troubleshooting SMF Services 3-60
- Debugging a Service That is Not Starting 3-61
- Restoring a Service in the Maintenance State 3-63
- Restoring a Service in the Maintenance State: Example 3-64
- Reverting to an SMF Snapshot 3-65
- Reverting to an SMF Snapshot: Example 3-66
- Repairing a Corrupt Repository 3-67
- Repairing a Corrupt Repository: Example 3-71
- Debugging Services During a System Boot 3-73
- Addressing system/filesystem/local:default Service Failures During Boot 3-74
- Summary 3-76

4 Administering ZFS

- Objectives 4-3
- Lesson Agenda 4-4
- Overview of ZFS 4-5
- Transactional File System 4-6
- Scalability 4-7
- Pooled Storage 4-8
- Dynamic Striping in a Storage Pool 4-10
- Data Integrity 4-11
- Mirrored Storage Pool Configuration (RAID-1) 4-12
- Parity Storage Pool Configuration (RAID-Z) 4-13
- ZFS File System 4-14

Snapshots	4-15
ZFS Clones	4-16
Quiz	4-17
Lesson Agenda	4-18
Administering ZFS Storage Pools	4-19
Creating ZFS Storage Pools	4-20
Determining Local Storage Disk Availability	4-21
Default Mount Point for Storage Pools	4-22
Creating a Basic ZFS Storage Pool	4-23
Creating a Mirrored Storage Pool	4-24
Creating a RAID-Z Storage Pool	4-25
Creating a ZFS Storage Pool with Log Devices	4-26
Creating a ZFS Storage Pool with Cache Devices	4-27
Displaying ZFS Storage Pool Information	4-28
Destroying ZFS Storage Pools	4-29
Managing ZFS Storage Pool Properties	4-30
Quiz	4-31
Lesson Agenda	4-32
Managing Devices in ZFS Storage Pools	4-33
Adding Devices to a Storage Pool	4-34
Attaching Devices to a Storage Pool	4-35
Detaching Devices from a Storage Pool	4-36
Taking Devices Offline in a Storage Pool	4-37
Bringing Devices Online in a Storage Pool	4-38
Replacing Devices in a Storage Pool	4-39
Designating Hot Spares in a Storage Pool	4-40
Creating Hot Spares in a Storage Pool	4-41
Adding Hot Spares to a Storage Pool	4-42
Replacing a Faulted Device With a Hot Spare	4-43
Removing Hot Spares in a Storage Pool	4-44
Quiz	4-45
Lesson Agenda	4-46
Administering ZFS File Systems	4-47
Creating a ZFS File System	4-48
Renaming a ZFS File System	4-49
Destroying a ZFS File System	4-50
Mounting ZFS File Systems	4-51
Unmounting a ZFS File System	4-53
ZFS File System Properties	4-54
ZFS File System Native Properties	4-55
Setting Quotas for ZFS File Systems	4-57

Setting Quotas for Users	4-58
Displaying User Space Usage	4-59
Removing User Quotas	4-60
Quiz	4-61
Lesson Agenda	4-62
Administering ZFS Snapshots and Clones	4-63
Creating a ZFS Snapshot	4-64
Displaying a ZFS Snapshot	4-65
Viewing Snapshot Space Accounting	4-66
Destroying a ZFS Snapshot	4-67
Renaming a ZFS Snapshot	4-68
Rolling Back a ZFS Snapshot	4-69
Identifying ZFS Snapshot Differences	4-70
Sending ZFS Snapshot Data	4-71
Receiving ZFS Snapshot Data	4-72
Replicating ZFS Snapshot Data Remotely	4-73
Creating a ZFS Clone	4-74
Destroying a ZFS Clone	4-75
Replacing a ZFS File System with a ZFS Clone	4-76
Quiz	4-78
Lesson Agenda	4-79
Securing ZFS File Systems	4-80
Delegated Administration	4-81
Delegating ZFS Permissions	4-82
Disabling ZFS Delegated Permissions	4-84
Removing ZFS Delegated Permissions	4-85
Data Encryption	4-86
Encrypting a ZFS Storage Pool and ZFS File System	4-87
Summary	4-88

5 Configuring the Network

Objectives	5-3
Agenda	5-4
Networking in Oracle Solaris 11	5-5
Network Stack in Oracle Solaris 11	5-6
Agenda	5-7
Prerequisites for Configuring a Network	5-8
Profile-Based Network Configuration	5-9
NCPs	5-10
Fixed NCP	5-11
Reactive NCP	5-12

Comparison Between Fixed and Reactive NCPs	5-13
Network Configuration and Administration Commands	5-14
netcfg Command	5-15
netadm Command	5-16
Commands to Configure Profiles	5-17
Configuring and Administering Datalink and Network Interfaces	5-18
dladm Command	5-19
dladm Types/Classes	5-20
Administering Datalinks with dladm Commands	5-21
ipadm Command	5-22
Administering Network Interfaces with the ipadm Command	5-23
Quiz	5-25
Agenda	5-26
Network Virtualization	5-27
Components of a Virtual Network	5-28
Network Virtualization in Zones	5-29
Network Virtualization in LDOMs	5-30
Configuring and Administering Virtual Networks	5-31
Creating a Virtual Network	5-32
Administering Virtual Networks	5-33
Migrating a VNIC	5-34
Quiz	5-35
Private Virtual Network	5-36
Features of a Private Virtual Network	5-37
Creating a Private Virtual Network	5-38
Establishing Communication Between Networks	5-39
Quiz	5-40
Agenda	5-41
High Availability	5-42
Overview of IPMP	5-43
IPMP Components	5-45
Types of IPMP Configurations	5-46
Failure and Repair Detection in IPMP	5-47
Configuring and Administering an IPMP Group	5-48
Creating an IPMP Group	5-49
Commands to Administer an IPMP Group	5-50
Quiz	5-51
Overview of Link Aggregation	5-52
Link Aggregation Types	5-53
Trunk Aggregation	5-54
LACP for Trunk Aggregation	5-55

Policies in Trunk Aggregation	5-56
Back-to-Back Configuration in Trunk Aggregation	5-57
Quiz	5-58
DLMP Aggregation	5-59
DLMP at Work	5-60
Comparison Between Trunk Aggregation and DLMP Aggregation	5-61
Preconfiguration Requirements for Link Aggregation	5-62
Creating a Link Aggregation	5-64
Commands to Administer Link Aggregations	5-65
Quiz	5-66
Agenda	5-67
Overview of Network Resource Management	5-68
Datalink Properties	5-69
Flows	5-70
Commands for Network Resource Management	5-71
dladm for Allocating Datalink Properties	5-72
flowadm for Managing Flows	5-73
Quiz	5-74
Managing Network Resources	5-75
Configuring Virtual Speed	5-76
Configuring CPU Pools for Datalinks	5-77
Allocating CPUs to Datalinks	5-78
Agenda	5-80
Need for Network Security	5-81
Overview of Link Protection	5-82
Link Protection Types	5-83
Configuring and Administering Link Protection	5-84
Summary	5-86

6 Administering Oracle Solaris Zones

Objectives	6-3
Lesson Agenda	6-4
Oracle Solaris Zones: Overview	6-5
Types of Zones	6-6
Zone States	6-7
Zone Commands	6-9
Quiz	6-10
Lesson Agenda	6-11
Zone Configuration Process	6-12
Creating a ZFS File System for Zones	6-14
Configuring a Zone	6-15

Displaying a Zone Configuration	6-17
Verifying That a Zone Is in the configured State	6-19
Gathering Information for the System Configuration Profile	6-20
Creating the SC Profile	6-21
Installing the Zone	6-22
Booting the Zone	6-23
Quiz	6-24
Lesson Agenda	6-25
Network Connectivity in Zones	6-26
Virtual Network Configuration	6-27
Checking the Virtual Network Configuration in a Zone	6-28
Verifying That a Zone's Virtual Network Interface Connection Is Operational	6-29
Quiz	6-30
Lesson Agenda	6-31
Administering an Oracle Solaris Zone	6-32
Displaying Zone Configuration Information	6-33
Logging In and Logging Out of a Zone	6-35
Halting, Shutting, and Starting a Zone	6-36
Quiz	6-37
Lesson Agenda	6-38
Zone Resource Management	6-39
Resource Pools	6-40
How Resource Pools Work	6-41
Allocating a Resource Pool to a Zone	6-42
Enabling Services for Resource Pools	6-43
Configuring a Persistent Resource Pool	6-44
Displaying the Resource Pool Configuration File	6-45
Modifying the Resource Pool Configuration File	6-47
Displaying and Committing the Modified Resource Pool Configuration File	6-49
Displaying the Resource Pool Configuration That Is Currently in Use	6-52
Displaying All Active Resource Pools	6-53
Binding the Zone to a Persistent Resource Pool	6-55
Allocating the Pool to the Zone	6-56
Rebooting the Zone	6-57
Confirming the Availability of the Resource Pool	6-58
Removing the Resource Pool Configuration	6-60
Removing the Resource Pool Configuration from the Zone	6-61
Rebooting the Zone	6-62
Checking the Resource Pool Configuration for the Zone	6-63
Deleting the Resource Pool	6-65
Resource Capping	6-66

Allocating Physical Memory Resources with Resource Capping	6-67
Quiz	6-68
Lesson Agenda	6-69
Securing Oracle Solaris Zones	6-70
Delegated Administration	6-71
Zone Link Protection	6-72
Exclusive IP	6-73
Immutable Zones	6-74
Cryptographic Services in Zones	6-75
Privileges	6-76
Users and Rights Profiles in Oracle Solaris Zones	6-78
Summary	6-79

7 Administering Privileges and RBAC

Objectives	7-3
Lesson Agenda	7-4
Assignment of User Privileges and Roles	7-5
Process Rights Management	7-6
Areas of Privilege	7-7
Sets of Privileges	7-8
Administering Privileges	7-10
Determining the Privileges Available to the Shell	7-11
Determining the Privileges on a Process	7-13
Displaying the Description of a Privilege	7-14
Determining the Privileges That Are Directly Assigned to You	7-15
Determining the Privileged Commands That a User Can Use	7-16
Assigning Privileges to a User or Role	7-17
Limiting the Privileges of a User or Role	7-18
Debugging Privilege Failure	7-19
Debugging Privilege Use in a Profile Shell	7-20
Debugging Privilege Use in a Regular Shell	7-21
Quiz	7-22
Lesson Agenda	7-23
Role-Based Access Control (RBAC)	7-24
Roles	7-25
Rights Profile	7-26
Authorizations	7-27
Privileges	7-28
Security Attributes	7-29
Key RBAC Files	7-30
Key RBAC Files: user_attr	7-31

- Key RBAC Files: auth_attr 7-32
- Key RBAC Files: exec_attr 7-34
- Key RBAC Files: prof_attr 7-36
- Relationships Among the RBAC Files 7-38
- Profile Shells 7-40
- Quiz 7-41
- Configuring RBAC 7-42
- Creating a Role 7-43
- Creating a Rights Profile 7-45
- Cloning a Rights Profile 7-46
- Modifying a Rights Profile 7-47
- Assigning a Rights Profile to a Role 7-49
- Assigning a Role to a User 7-50
- Assuming a Role 7-52
- Restricting an Administrator to Explicitly Assigned Rights 7-53
- Assigning the Rights Profile to a User 7-54
- Delegating an Authorization to a User 7-55
- Assigning Authorization to a Role 7-57
- Rights Profiles 7-58
- Modifying a System-Wide RBAC Policy 7-59
- Quiz 7-61
- Summary 7-62

8 Installing the Oracle Solaris 11 Operating System

- Objectives 8-3
- Lesson Agenda 8-4
- Oracle Solaris Installation 8-5
- Preparing for the Installation 8-6
- Reviewing the Release Notes 8-7
- Selecting the Installation Option 8-8
- Identifying System Requirements 8-9
- Downloading Images 8-10
- Agenda 8-11
- Installing Oracle Solaris 11 by Using the Live Media Installer 8-12
- Introducing the Live Media Desktop 8-14
- Initiating the Installation with Live Media 8-15
- Welcome Screen 8-16
- Oracle Solaris 11 Live Media: Disk Discovery 8-17
- Selecting a Disk 8-18
- Setting the Time Zone, Date, and Time 8-19
- Providing User Information 8-20

Support Registration	8-21
Reviewing Installation Specifications	8-22
Monitoring the Installation	8-23
Verifying the Installation	8-24
Reviewing the Installation Log	8-25
Rebooting the System	8-28
Login Screen	8-29
Checking the Login Username	8-30
Checking the Login Password	8-31
Lesson Agenda	8-32
Installing Oracle Solaris 11 by Using the Text Installer	8-33
Verifying the Installation	8-34
Verifying Login Information	8-35
Verifying the System's Host Name	8-36
Displaying Basic System Information	8-37
Displaying the System's Release Information	8-38
Displaying Disk Configuration Information	8-39
Displaying Disk Configuration Information: Format Menu	8-40
Displaying Disk Configuration Information: Partition Table	8-41
Displaying the Installed Memory Size	8-42
Displaying Information About Network Services	8-43
Displaying Network Interface Information	8-44
Baseline System Information Commands: Summary	8-45
Quiz	8-46
Lesson Agenda	8-47
Automated Installer (AI): Overview	8-48
Automated Installer: Components	8-49
Automated Installer: Process	8-50
Automated Installer: Flowchart	8-51
Performing an AI Installation	8-52
Reviewing AI Installation Server Requirements	8-53
Verifying the Server Software Requirements	8-54
Verifying the Static IP Address	8-55
Verifying That DNS Is Operational	8-56
Enabling the DNS Multicast Service	8-57
Verifying That IPS Is Available Locally	8-58
Verifying That the DHCP Server Is Enabled	8-59
Configuring the AI Server	8-60
Installing the OS on the Client System	8-61
Identifying Client System Requirements	8-62
Identifying the Installation Files	8-63

Performing the Installation	8-64
Reviewing Client Installation Messages	8-65
Quiz	8-67
Summary	8-68

9 Monitoring System Resources


Objectives	9-3
Agenda	9-4
Monitoring and Observability Tools	9-5
iostat Utility	9-6
kstat Utility	9-7
mpstat Utility	9-8
pgstat Utility	9-9
fsstat Utility	9-10
poolstat Utility	9-11
svcs Utility	9-12
netstat Utility	9-13
dlstat Utility	9-14
flowstat Utility	9-15
ipmpstat Utility	9-16
acctadm Utility	9-17
zonestat Utility	9-18
vmstat Utility	9-19
prstat Utility	9-20
truss Utility	9-21
ptree Utility	9-22
Quiz	9-23
Agenda	9-24
DTrace: Overview	9-25
DTrace: Capabilities	9-26
DTrace: Components	9-27
Probes	9-28
Providers	9-29
Consumers	9-30
D Language	9-31
DTrace Toolkit	9-32
DTrace Toolkit: Important Scripts	9-33
Before Using DTrace	9-34
Launching DTrace	9-35
DTrace: Example	9-36
Summary	9-37

6

Administering Oracle Solaris Zones

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.



System Administration for Experienced UNIX/Linux Administrators



Administering System
Software by Using IPS



Administering Services
by Using SMF



Administering ZFS



Configuring the Network



Administering Oracle Solaris
Zones



Administering Privileges
and RBAC



Installing the Oracle Solaris 11
Operating System



Monitoring System Resources

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Objectives

After completing this lesson, you should be able to:

- Explain the fundamentals of Oracle Solaris Zones
- Configure Oracle Solaris Zones
- Configure network connectivity in Oracle Solaris Zones
- Administer Oracle Solaris Zones
- Manage system resources in Oracle Solaris Zones
- Secure Oracle Solaris Zones

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Lesson Agenda

- Explaining the fundamentals of Oracle Solaris Zones
- Configuring Oracle Solaris Zones
- Configuring network connectivity in Oracle Solaris Zones
- Administering Oracle Solaris Zones
- Managing system resources in Oracle Solaris Zones
- Securing Oracle Solaris Zones

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Oracle Solaris Zones: Overview

- Oracle Solaris Zones are virtualized operating system (OS) environments. Each zone is created within a single instance of the Oracle Solaris OS.
- A zone is an application execution environment in which processes are isolated from the rest of the system.
- This isolation prevents processes that are running in one zone from monitoring or affecting processes that are running in other zones.
- Even a process running with superuser credentials cannot view or affect activity in other zones.

The Oracle logo, consisting of the word "ORACLE" in white, uppercase letters on a red rectangular background.

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Oracle Solaris Zones:

- Provide the opportunity to consolidate numerous applications onto fewer, more scalable servers
- Allow applications to be consolidated with other applications
- Allow dynamic resource reallocation, which permits unused resources to be shifted to other zones as needed
- Provide an isolated and secure environment for running applications
- Enable a one-application-per-server deployment model while simultaneously sharing hardware resources
- Can be installed and run on shared storage

Note: With Oracle Solaris 11.1, you can configure, install, and run Oracle Solaris Zones hosted directly on arbitrary storage device objects, such as Fibre Channel or iSCSI targets.

Types of Zones

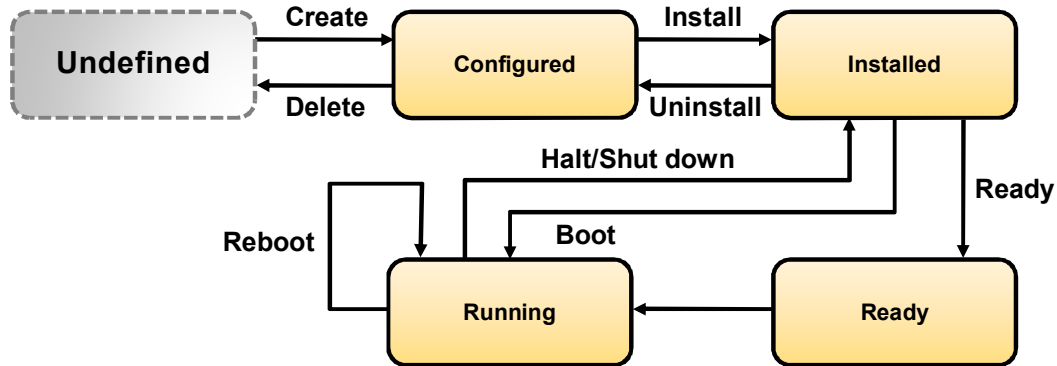
- Zones can be used on any machine that is running Oracle Solaris 10 or later versions.
- The number of zones is determined by the:
 - Total resource requirements of the application software running in all the zones
 - Size of the system
- Oracle Solaris supports the following types of zones:
 - Global zone
 - Non-global zone
 - Branded zone

The Oracle logo, consisting of the word "ORACLE" in a white, sans-serif font, centered within a red rectangular background.

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

- **Global zone:** Every Oracle Solaris system contains a global zone. The global zone has a dual function. It is the default zone for the system and is used for system-wide administrative control. Only the global zone is bootable from the system hardware. The global zone is the only zone from which a non-global zone can be configured, installed, managed, or uninstalled. Note that all applications run in the global zone if no non-global zones are created.
- **Non-global zones:** Simply called *zones*, are configured inside the global zone. A zone provides isolation at almost any level of granularity you require. A zone can be thought of as a box in which one or more applications can be run without interacting with the rest of the system. A zone does not need a dedicated CPU, a physical device, or a portion of physical memory. These resources can either be multiplexed across a number of zones running within a system, or allocated on a per-zone basis using the resource management features available in the operating system.
- **Branded zones (BrandZ):** Is an extension of Oracle Solaris Zones. The BrandZ framework is used to create non-global branded zones that contain operating environments that are different from that of the global zone. For example, you can run Oracle Solaris 10 applications by using Oracle Solaris 10 Zones (`solaris10` brand) on a system running Oracle Solaris 11.

Zone States



ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

As a zone is configured, enabled, and used, its status changes.

Non-Global Zone States

- **Undefined:** In this state, the zone's configuration has not been completed and committed to stable storage. This state also occurs when a zone's configuration has been deleted.
- **Configured:** In this state, the zone's configuration is complete and committed to stable storage. However, those elements of the zone's application environment that must be specified after initial boot are not yet present.
- **Incomplete:** This is a transitional state. During an install or uninstall operation, the state of the target zone is set to incomplete. After successful completion of the operation, the state is set to the correct state. However, a zone that is unable to complete the install process will stop in this state.
- **Unavailable:** The zone enters this state when the zone's storage is unavailable, the zone's software is incompatible with the global zone's software, or an archive is successfully extracted but the installations fail. This state allows `pkg` operations to work even if a zone's storage is not available.

- **Installed:** In this state, the zone configuration is instantiated on the system. At this point, the system administrator verifies that the configuration can be successfully used on the designated Oracle Solaris system. Packages are installed under the zone's root path. In this state, the zone has no associated virtual platform.
- **Ready:** In this state, the virtual platform for the zone is established. The kernel creates the zone scheduling process, network interfaces are set up and made available to the zone, file systems are mounted, and devices are configured. A unique zone ID is assigned by the system. At this stage, no processes associated with the zone have been started.
- **Running:** In this state, the user processes associated with the zone application environment are running. The zone enters the running state as soon as the first user process associated with the application environment (`init`) is created.
- **Shutting down, down:** These states are transitional states that are visible while the zone is being halted. However, a zone that is unable to shut down for any reason will stop in one of these states.

Zone Commands

The commands summarized in the following table provide the primary administrative interface to the zones facility.

Command	Description
zlogin	Log in to a non-global zone.
zoneadm	Administer zones on a system.
zonecfg	Perform zone configuration.
zonestat	Observe and monitor zone resource usage.

The Oracle logo, consisting of the word "ORACLE" in a white, sans-serif font, is positioned on the right side of a solid red horizontal bar.

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Quiz

In the process of consolidation, Joe wants to migrate legacy data residing on an Oracle Solaris 10 system to an Oracle Solaris 11 system. At the same time, he needs workload separation.

What type of zone should Joe create in the host operating system to maintain his Oracle Solaris 10 data?

- a. Non-global zone
- b. Global zone
- c. Branded zone

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Answer: c

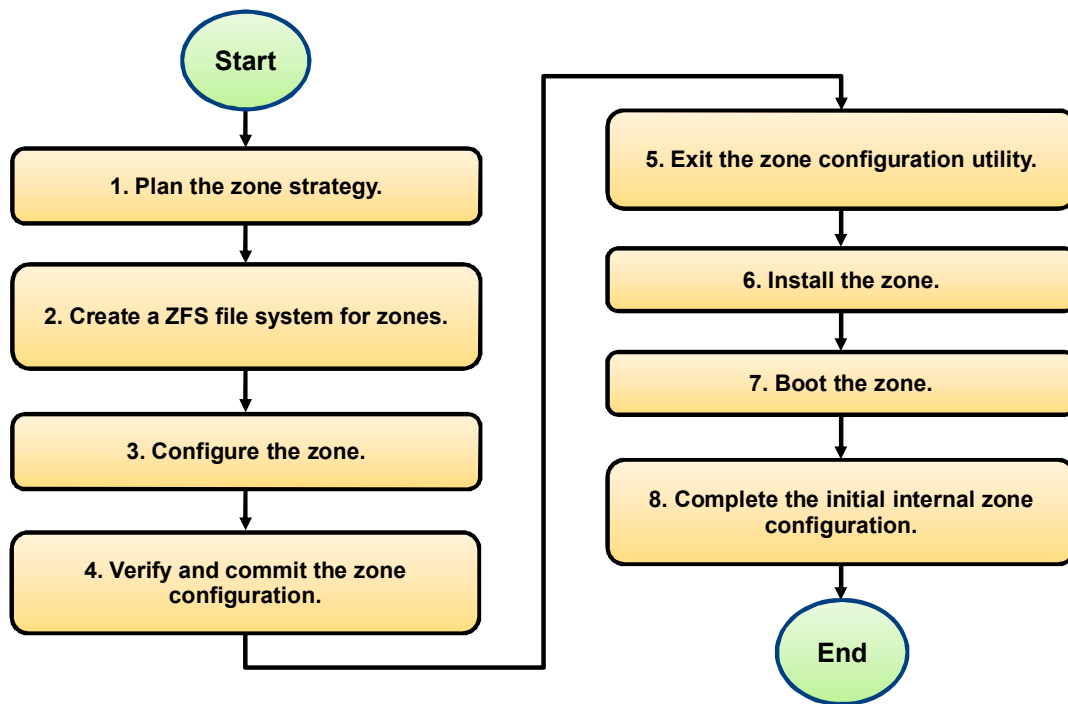
Lesson Agenda

- Explaining the fundamentals of Oracle Solaris Zones
- **Configuring Oracle Solaris Zones**
- Configuring network connectivity in Oracle Solaris Zones
- Administering Oracle Solaris Zones
- Managing system resources in Oracle Solaris Zones
- Securing Oracle Solaris Zones

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Zone Configuration Process



ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Configuring a zone involves the following general steps:

1. **Planning the zone strategy:** This is a critical step to ensure that you accurately understand your zone requirements in terms of the number of zones, the zone names, the zone paths, the IP types, and so on.
2. **Creating the ZFS file system for all the zones:** All the required system software and any additional packages are installed in the private file systems of the zone.
3. **Configuring the zone:** During configuration, the administrator identifies the resources and properties for the zone and identifies whether the zone is to be a shared-IP zone or an exclusive-IP zone.
4. **Verifying and committing the zone configuration:** The verification step ensures that the configuration of the specified zone can be safely installed on the system. The commit command takes the configuration from memory and puts it into permanent storage.
5. **Exiting the zone configuration utility:** The administrator exits the zone configuration utility.

6. **Installing the zone:** The install process automatically creates a ZFS file system (dataset) for the zone path when the zone is installed. If a ZFS dataset cannot be created, the zone is not installed. The zone installation packages are installed from the Image Packaging System (IPS).
Note: Oracle Solaris zones participate fully in the boot environment (BE) framework. When a system is upgraded by using IPS, all the zones on that system are cloned by using ZFS cloning. On system reboot, the newly updated clones of the zones become active. This ensures both a clean upgrade process and the ability to roll back if something fails.
7. **Booting the zone:** After the configured zone has been installed, it can be booted or activated.
8. **Completing initial internal zone configuration:** The internal configuration specifies a naming service to use, the default locale and time zone, the zone's root password, and other aspects of the OS environment. The zone is now ready to be used.

Creating a ZFS File System for Zones

- To create a file system for a zone, use the `zfs create` command.

```
# zpool create zonepool mirror c0t0do c1t1d1  
# zfs create -o mountpoint=/zones zonepool/zones
```

- To verify that the file system exists and that it has been mounted, use the `zfs list` command.

```
# zfs list zonepool/zones  
NAME                USED  AVAIL  REFER  MOUNTPOINT  
rpool/zones         31K   22.6G   31K    /zones
```



Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

To create a file system for a zone, use the `zfs create` command with the `-o` option (to specify the `mountpoint` property), followed by the `mountpoint` property value (`mountpoint=/zones`) and the file system name (`zonepool/zones`), as shown in the first example in the slide.

You can then use the `zfs list` command followed by the file system name to verify that the file system has been created and mounted, as shown in the second example.

Configuring a Zone

To configure a zone, use `zonecfg -z zonename`.

```
# zonecfg -z hrzone
Use 'create' to begin configuring a new zone.
zonecfg:hrzone> create
create: Using system default template 'SYSdefault'
zonecfg:hrzone> set zonepath=/zones/hrzone
zonecfg:hrzone> set autoboot=true
zonecfg:hrzone> add net
zonecfg:hrzone:net> set physical=vnic1
zonecfg:hrzone:net> end
zonecfg:hrzone> verify
zonecfg:hrzone> commit
zonecfg:hrzone> exit
```

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

The `zonecfg` command is used to create the zone configuration. You must be a superuser or have the appropriate rights profile to configure a zone.

To perform the configuration:

1. Use the `zonecfg` command with the `-z` option followed by the zone name to specify the name of the zone, as shown in the example. After you enter the command (if you are configuring a new zone), you see the following message: Use 'create' to begin configuring a new zone.
2. At the zone prompt, enter `create`. This enables you to create the new zone configuration by setting specific properties, such as the zone path, the IP type, and the network type.
3. Set the zone path by using the `set zonepath=` command followed by the zone name (for example, `/zones/hrzone`).
4. Set `autoboot` to `true` by using `set autoboot=true`. This setting indicates that the zone should be booted automatically at system boot.
5. To add a network interface to the zone, use the `add net` command. Notice that the `zonecfg` prompt changes its scope to include "net": `zonecfg:hrzone:net`, where you can set network resource properties for the zone. For example, `set physical=vnic1` sets the `physical` property to `vnic1` for `hrzone`. Adding other resources, such as memory and file system, follows the same pattern.

6. Exit the scope by using the `end` command.
7. After you complete your zone configuration, you need to verify that all required information is present. You do this by using the `verify` command.
 - If all required information is not present, the system notifies you. You must then review your configuration to determine what is missing.
 - If no messages are displayed, you can continue to the next step.
8. The `commit` command takes the zone configuration from memory and puts it into permanent storage.
9. After the zone configuration is committed, you can exit the zone configuration session by using the `exit` command.

Displaying a Zone Configuration

To display a zone configuration, use `zonecfg -z zonename info`.

```
# zonecfg -z hrzone info
zonename: hrzone
zonepath: /zones/hrzone
brand: solaris
autoboot: true
file-mac-profile:
bootargs:
pool:
limitpriv:
scheduling-class:
ip-type: exclusive
hostid:
fs-allowed:
net:
    address not specified
    allowed-address not specified
    physical: vnic1
    defrouter not specified
<continued on next slide>
```

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

After finishing your zone configuration, it is a good practice to review it before you install the zone. To display a zone configuration, use the `zonecfg -z` command followed by the zone name and the `info` subcommand, as shown in the slide. Verify that you have correctly set the zone path, IP type, and network interface properties.

Displaying a Zone Configuration

```
<continued from previous slide>
anet:
  linkname: net0
  lower-link: auto
  allowed-address not specified
  configure-allowed-address: true
  defrouter not specified
  allowed-dhcp-cids not specified
  link-protection: mac-nospoof
  mac-address: random
  mac-prefix not specified
  mac-slot not specified
  vlan-id not specified
  priority not specified
  rxrings not specified
  txrings not specified
  mtu not specified
  maxbw not specified
  rxfanout not specified
  vsi-typeid not specified
  vsi-vers not specified
  vsi-mgrid not specified
  etsbw-lcl not specified
  cos not specified
  pkey not specified
  linkmode not specified
```

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Note that some properties such as `linkname` and `mac-address` are automatically picked from the default zone template found in the `/etc/zones` directory. You can remove such default properties by using the `zonecfg` command. For example, to remove `net0`, use the following command:

```
# zonecfg -z zone1 'remove anet linkname=net0'
```

Verifying That a Zone Is in the configured State

To list all configured and running zones on the system, use `zoneadm list -cv`.

# zoneadm list -cv					
ID	NAME	STATUS	PATH	BRAND	IP
0	global	running	/	solaris	shared
-	hrzone	configured	/zones/hrzone	solaris	excl
-	itzone	configured	/zones/itzone	solaris	excl

The Oracle logo, consisting of the word "ORACLE" in a white, sans-serif font, centered on a red rectangular background.

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

You are now ready to install the zone. But it is a good practice to first confirm that the zone is in the configured state. You can use the `zoneadm list -cv` command to see all configured and running zones on a system, as shown in the example in the slide.

Gathering Information for the System Configuration Profile

- After verifying that the zone is in the `configured` state, you must create a system configuration (SC) profile for the zone.
- The SC profile specifies the default locale and time zone, the zone's `root` password, a naming service to use, and other aspects of the application environment, including (but not limited to) the following:
 - Computer name of the zone (for example, `hrzone`)
 - IP address of the zone, which is based on the IP address of the zone's VNIC
 - Netmask of the IP address

The Oracle logo, consisting of the word "ORACLE" in a white, sans-serif font, centered on a red rectangular background.

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

This slide presents a sample of the type of information that you need to supply before you can complete the system configuration profile.

- Computer name: `hrzone`
- Wired Ethernet network configuration: `Manually`
- IP address of the zone: `192.168.1.100`
- Netmask of the IP address: `255.255.255.0`
- DNS name service: `Do not configure DNS`
- Alternate name service: `None`
- Time zone, region, and location: *Use your local region.*
- Users, username, and password

You must gather this information before creating the SC profile. Most of the information is supplied by selecting from a list of choices. Typically, the default options are enough unless your system configuration requires something else. After you supply the required information for the zone, the zone is restarted.

Creating the SC Profile

To create the SC profile, use `sysconfig create-profile -o pathname`.

```
# sysconfig create-profile -o /opt/ora/data/hrconf.xml  
  
<prompt sequence omitted>  
  
Exiting System Configuration Tool. Log is available at:  
/system/volatile/sysconfig/sysconfig.log.XXXX
```

The Oracle logo, consisting of the word "ORACLE" in a white, sans-serif font, is centered on a solid red rectangular background.

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

To create a SC profile for a zone, use the `sysconfig create-profile -o` command followed by the path name of the location in which you want the profile to reside, as shown in the example in the slide. Using the configuration information that you gathered previously, you need to respond to each of the prompts. When you finish, you are exited from the System Configuration Tool.

Installing the Zone

To install a zone, use `zoneadm -z zonename install -c profile_pathname`.

```
# zoneadm -z hrzone install -c /opt/ora/data/hrconf.xml
The following ZFS file system(s) have been created:
    rpool/zones/hrzone
Progress being logged to
/var/log/zones/zoneadm.20131119T021542Z.hrzone.install
.....
    Done: Installation completed in 188.189 seconds.
    Next Steps: Boot the zone, then log into the zone console
(zlogin -C) to complete the configuration process.
Log saved in non-global zone as
/zone/hrzone/root/var/log/zones/zoneadm.20131119T021542Z.hrzone
.install
#
```



Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

After you have created the SC profile, you are ready to install the zone. To install a zone, use the `zoneadm -z` command followed by the zone name, the `install -c` subcommand, and the path name to the SC profile, as shown in the example in the slide.

The installation process automatically creates a ZFS file system (dataset) for the zone path when the zone is installed. If the file system cannot be created, the zone is not installed. The installation process also verifies the specified publisher and downloads the zone installation packages from IPS. This process normally takes about three to five minutes for each zone (depending on your system).

Booting the Zone

- To list all running and installed zones on the system, use `zoneadm list -iv`.

```
# zoneadm list -iv
```

ID	NAME	STATUS	PATH	BRAND	IP
0	global	running	/	solaris	shared
-	hrzone	installed	/zones/hrzone	solaris	excl
-	itzone	installed	/zones/itzone	solaris	excl

- To boot a zone, use `zoneadm -z zonename boot`.

```
# zoneadm -z hrzone boot
# zoneadm -z itzone boot
# zoneadm list -v
```

ID	NAME	STATUS	PATH	BRAND	IP
0	global	running	/	solaris	shared
1	hrzone	running	/zones/hrzone	solaris	excl
2	itzone	running	/zones/itzone	solaris	excl

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

The next step is to boot the zone. But it is a good practice to first confirm that the zone is in the installed state. You can use the `zoneadm list -iv` command to see all the running and installed zones on a system, as shown in the first example in the slide. Observe that both `hrzone` and `itzone` have a status of installed.

You can now boot the installed zones. To boot a zone, use the `zoneadm -z` command followed by the zone name and the `boot` subcommand, as shown in the second example.

To verify that a zone is in the running state, you can run the `zoneadm list -v` command, as shown in the second part of the second example. Note that the two non-global zones now have assigned IDs.

Quiz

Which of the following is an incorrect command in the sequence of configuring a non-global zone?

- a. Create a ZFS file system for zones.
(`zfs create file_system`)
- b. Remove the `net0` interface.
(`zonecfg -z zonename 'remove anet linkname=net0'`)
- c. Configure the zone.
(`zonecfg -z zonename`)
- d. Install the zone.
(`zoneadm -z zonename install -c profile_pathname`)
- e. Boot the zone.
(`zoneadm -z zonename boot`)

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Answer: b

Lesson Agenda

- Explaining the fundamentals of Oracle Solaris Zones
- Configuring Oracle Solaris Zones
- **Configuring network connectivity in Oracle Solaris Zones**
- Administering Oracle Solaris Zones
- Managing system resources in Oracle Solaris Zones
- Securing Oracle Solaris Zones

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Network Connectivity in Zones

- Zones communicate through IP network interfaces.
- The system administrator configures zone network interfaces during zone configuration.
- When a zone is booted, the network interfaces are set up and placed in the zone.
- Two IP types are available for non-global zones:
 - **Shared-IP:** The network interface is shared with the global zone.
 - **Exclusive-IP:** The network interface is dedicated to the non-global zone.

The Oracle logo, consisting of the word "ORACLE" in a white, sans-serif font, is centered on a solid red rectangular background.

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Basic communication between zones is accomplished by giving each zone Internet Protocol (IP) network connectivity. An application running in one zone cannot observe the network traffic of another zone. This isolation is maintained even though the respective streams of packets travel through the same physical interface.


During zone configuration, the system administrator configures the zone network interfaces to provide network connectivity. These network interfaces are set up and placed in the zone when it is booted.

Two IP types are available for non-global zones:

- **Shared-IP:** Zones share a network interface (or IP layer) with the global zone.
- **Exclusive-IP:** Zones have their own dedicated network interface (or instance of the IP layer).

Note: Both shared-IP and exclusive-IP can be used in a system. However, exclusive-IP is the default.

Virtual Network Configuration



ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

The graphic in the slide illustrates a simple virtual network configuration. There are two zones, `hrzone` and `itzone`, each with a dedicated or exclusive IP address. The `hrzone` zone uses `vnic1` as its network interface, and `itzone` uses `vnic2` as its network interface.

Checking the Virtual Network Configuration in a Zone

To display the network interface address information for a zone, log in to the zone and then use `ipadm show-addr`.

```
# zlogin hrzone
[Connected to zone 'hrzone' pts/2]
Oracle Corporation      SunOS 5.11      11.0      September 2012
root@hrzone:~# ipadm show-addr
```

ADDROBJ	TYPE	STATE	ADDR
lo0/v4	static	ok	127.0.0.1/8
vnic1/v4	static	ok	192.168.1.100/24
lo0/v6	static	ok	::1/128
vnic1/v6	addrconf	ok	fe80::8:20ff:fe43:7986/10

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

To display the virtual network configuration in a zone, you first need to log in to the zone by using the `zlogin` command followed by the zone name, as shown in the example in the slide. After you are logged in, you can use the `ipadm show-addr` command to see the network interface address information for the zone.

Verifying That a Zone's Virtual Network Interface Connection Is Operational

Use `ping` and the IP address.

```
root@hrzone:~# ping 192.168.1.200
192.168.1.200 is alive
```

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

To verify that a zone's network connection is operational, ping the IP address from within the zone.

Quiz

Joe has a BrandZ and a non-global zone installed on his system. He intends to make `ip-type=shared` for the branded zone and `ip-type=exclusive` for the non-global zone. Does Oracle Solaris 11.1 support both IP types in a system?

- a. Yes
- b. No

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Answer: a

Lesson Agenda

- Explaining the fundamentals of Oracle Solaris Zones
- Configuring Oracle Solaris Zones
- Configuring network connectivity in Oracle Solaris Zones
- **Administering Oracle Solaris Zones**
- Managing system resources in Oracle Solaris Zones
- Securing Oracle Solaris Zones

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Administering an Oracle Solaris Zone

Administering an Oracle Solaris Zone involves the following activities:

- Displaying zone configuration information
- Logging in and logging out of a zone
- Halting, shutting down, and starting a zone

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Displaying Zone Configuration Information

To display a zone configuration, use `zonecfg -z zonename info`.

```
# zonecfg -z hrzone info
zonename: test1
zonepath: /zones/hrzone
brand: solaris
autoboot: true
bootargs:
pool:
limitpriv:
scheduling-class:
ip-type: shared
hostid:
fs-allowed:
[max-lwps: 500]
<output continued on next slide...>
```

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

To display a zone's configuration, you must be the global administrator in the global zone or a user with the correct rights profile. To display the configuration, use the `zonecfg -z` command followed by the zone name and the `info` subcommand, as shown in this example.

The first part of the output displays the zone name (`hrzone`), the zone path (`/zones/hrzone`), the brand (`solaris`), and the setting of the `autoboot` option. The `autoboot` option, when set to `true`, indicates that the zone should be booted automatically at system boot.

Notice also the IP type setting. In this example, `ip-type` is `shared`, which means that this non-global zone shares the IP layer with the global zone.

Displaying Zone Configuration Information

```
<...output continued from previous slide>
fs:
  dir: /local/hrzone
  special: rpool/hrzone
  raw not specified
  type: lofs
  options: []
net:
  address: 192.168.0.200
  allowed-address not specified
  configure-allowed-address: true
  physical: net0
  defrouter not specified
rctl:
  name: zone.max-lwps
  value: (priv=privileged,limit=500,action=deny)
```

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

As shown in the slide, the next section of the output displays:

- Zone file system information (`fs`)
- Zone network information, such as the IP address (`192.168.0.200`) and NIC (`net0`)
- Zone attribute information
- Resource control settings (`rctl`)

Note: If a zone uses a virtual network interface, the VNIC interface is displayed after `physical` in the network section.

Logging In and Logging Out of a Zone

- To log in to a zone, use `zlogin` followed by the zone name.

```
# zlogin hrzone
[Connected to zone 'hrzone' pts/1]
Oracle Corporation SunOS 5.11      11.1   September 2012
```

- To exit a non-global zone from a pseudo terminal or terminal login, use `exit`.

```
# exit
```

- To disconnect from a zone from a virtual console or console login, use `~.`

```
# ~.
```

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

You must be logged in to a zone to perform administrative tasks in that zone, such as modifying the configuration, taking a backup, or monitoring resource usage.

The `zlogin` utility is used to enter a non-global zone (see the first example in the slide). The utility can be used only from the global zone as a root user or as a user with required privileges.

When you have completed your administrative tasks in a zone, you must log out of, or exit, the zone. Logging out of the zone can save on system resources, especially if you are running multiple zones on a system.

There are two ways to exit a non-global zone:

- From a non-virtual console:** Use the `exit` command, as shown in the second example.
- From a virtual console:** Use the tilde (`~`) character and a period (`.`), as shown in the third example.

Halting, Shutting, and Starting a Zone

- To halt a zone, use `zoneadm -z zonename halt`.

```
# zoneadm -z hrzone halt
```

- To shut down a zone, use `zoneadm -z zonename shutdown`.

```
# zoneadm -z hrzone shutdown
```

- To start a zone, use `zoneadm -z zonename boot`.

```
# zoneadm -z hrzone boot
```

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

When you need to remove a zone, use the `zoneadm -z halt` command with the zone name, as shown in the example. When halted, the zone is brought back to the `installed` state. To verify that the zone has been halted and is no longer running, you can run the `zoneadm list -iv` command.

Although you can halt a zone, the recommended way to bring a zone down is by using the `zoneadm shutdown` command. This approach brings the zone down more gently. To shut down a zone from the global zone, use the `zoneadm -z` command followed by the zone name and specify `shutdown` as the command to run, as shown in the second example. This procedure is used to cleanly shut down a zone (as opposed to halting it).

After a zone has been shut down, you can start it again by using the `zoneadm -z` command followed by the zone name and `boot`, as shown in the third example. When a zone is booted, the required services and facilities for that zone are brought online.

Quiz

The `zlogin` utility can be used only from a non-global zone by a user with required privileges.

- a. True
- b. False

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Answer: b

Lesson Agenda

- Explaining the fundamentals of Oracle Solaris Zones
- Configuring Oracle Solaris Zones
- Configuring network connectivity in Oracle Solaris Zones
- Administering Oracle Solaris Zones
- **Managing system resources in Oracle Solaris Zones**
- Securing Oracle Solaris Zones

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Zone Resource Management

- Resource management:
 - Uses a collection of algorithms to minimize cross-workload performance compromises
 - Provides facilities that enable you to modify the default behavior of the operating system with respect to different workloads
- Resources that can be managed in a zone include:
 - Resource pools or assigned CPUs
 - Resource capping

The Oracle logo, consisting of the word "ORACLE" in white capital letters on a red rectangular background.

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Oracle Solaris 11 provides the option to manage varying workloads that are generated by different applications on a system. A workload is an aggregation of all processes of an application or a group of applications.

If you choose not to use the resource management features, the Oracle Solaris OS responds to workload demands by adapting to new application requests dynamically. This default response generally means that all activity on the system is given equal access to resources. Oracle Solaris resource management features enable you to treat workloads individually.

If you use resource management features, you should align the boundaries of the resource management controls with those of the zones. This alignment creates a more complete model of a virtual machine, where namespace access, security isolation, and resource usage are all controlled. Resources that can be controlled in a zone include the following:

- **Resource pool or assigned CPU:** Is used for partitioning machine resources. A resource pool represents an association between groups of resources that can be partitioned.
- **Resource cap:** Is a per-process, per-task, per-project limit on the consumption of a resource. A *project* is a network-wide administrative identifier for related work.

Resource Pools

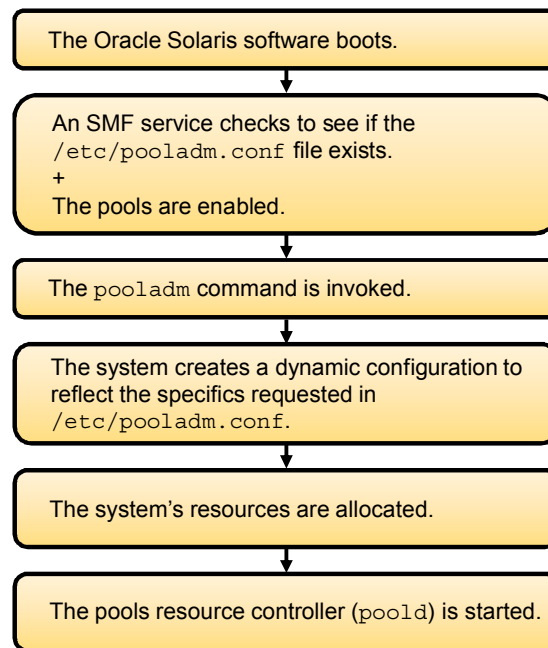
- Resource pool allocation enables you to separate workloads so that the consumption of certain resources does not overlap with the consumption of other resources.
- This resource reservation helps to achieve predictable performance on systems with mixed workloads.
- SMF supports two resource pool services:
 - Default resource pool service
`svc:/system/pools:default`
 - Dynamic resource pool service
`svc:/system/pools/dynamic:default`
- Resource pool services are disabled by default.

The Oracle logo, consisting of the word "ORACLE" in a white, sans-serif font, is positioned on the right side of a solid red horizontal bar.

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

The slide lists the two types of resource pool services in the Oracle Solaris service management facility (SMF) that reside on the system. Note that the dynamic pool service is dependent on the default pool service. By default, neither service is active.

How Resource Pools Work

**ORACLE**

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

When the Oracle Solaris software boots, an SMF service checks to see if the `/etc/pooladm.conf` file exists. If this file is found and the pools are enabled, the `pooladm` command is invoked to make this configuration the active pools' configuration. The system creates a dynamic configuration to reflect the specifics requested in `/etc/pooladm.conf`, and the machine's resources are allocated accordingly.

Note: The `pooladm` command is used to activate and deactivate the resource pools facility.

The pools resource controller, `poold`, is started with the dynamic resource pools facility. This system daemon should always be active when dynamic resource allocation is required. The `poold` resource controller identifies available resources and monitors workloads to determine when system usage objectives are no longer being met. The controller then considers alternative configurations in terms of the objectives, and remedial action is taken. If possible, the resources are reconfigured so that the objectives can be met. If this action is not possible, the daemon logs that the user-specified objectives can no longer be achieved. Following a reconfiguration, the daemon resumes monitoring workload objectives.

Allocating a Resource Pool to a Zone

1. Enable the two resource pool services.
2. Create a pool configuration file and save it in the default configuration file `/etc/pooladm.conf`.
3. Modify the pool configuration file to specify a subset of the system's processors that should be dedicated to a zone.
4. Bind the resource pool to the zone.

The Oracle logo, consisting of the word "ORACLE" in a white, sans-serif font, centered within a solid red rectangular background.

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

To allocate a resource pool to a zone, you must first enable the resource pool services and create a pool configuration file for the current pool configuration that you save in the default `/etc/pooladm.conf` configuration file. This file, which is in XML format, contains a description of the pools to be created on the system and the elements that can be manipulated: system, pool, pset (processor set), and cpu. This configuration file is often referred to as the *static configuration file*.

After creating and saving the configuration file, you can modify it to specify a subset of the system's processors that should be dedicated to a zone while it is running. The static configuration file now matches the current dynamic configuration that represents the way you want the system to be configured with respect to how the resource pool or pools will function. After modifying the static configuration file and saving the changes, you must allocate or bind the zone to the resource pool.

Enabling Services for Resource Pools

- To activate the resource pool services, use `svcadm enable -r pools/dynamic`.

```
# svcadm enable -r pools/dynamic
```

- To verify that the service pools and the `poold` daemon are up, use `svcs *pools*` and `pgrep -lf poold`, respectively.

```
# svcs *pools*
STATE          STIME          FMRI
online         16:08:10      svc:/system/pools:default
online         16:08:11      svc:/system/pools/dynamic:default
# pgrep -lf poold
2283 /usr/lib/pool/poold
```

The Oracle logo, consisting of the word "ORACLE" in a white, sans-serif font on a red rectangular background.

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Resource pool services are not enabled by default. To activate both resource pool services, you can run the `svcadm enable` command with the `-r` option, followed by `pools/dynamic`, as shown in the example in the slide.

To verify that the pool services are online, you can run the `svcs *pools*` command. You can also verify that the `poold` daemon is running.

Configuring a Persistent Resource Pool

- To create the pool configuration file, use `pooladm -s`.

```
# pooladm -s
```

- To verify that the file has been created, use `ls -l /etc/pool*`.

```
# ls -l /etc/pool*  
-rw-r--r-- 1 root root 1298 Dec 14 16:13 /etc/pooladm.conf  
# file /etc/pooladm.conf  
/etc/pooladm.conf:      XML document
```

The Oracle logo, consisting of the word "ORACLE" in a white, sans-serif font on a red rectangular background.

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

After enabling the resource pool services, you can create the pool configuration file and save it in the default configuration file `/etc/pooladm.conf`. To create the pool configuration file, use the `pooladm` command with the `-s` option, which saves the file.

Note: The `pooladm` command is used to activate and deactivate the resource pools facility. To verify that the file is created, you can use the `ls -l /etc/pool*` command, as shown in the example in the slide. If you run the `file /etc/pooladm.conf` command, you can see that the file is an XML document.

Displaying the Resource Pool Configuration File

Use `poolcfg -c info`.

```
# poolcfg -c info
system default
    string  system.comment
    int     system.version 1
    boolean system.bind-default true
    string  system.poolid.objectives wt-load

<Complete output presented in the Notes>
```

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Before you modify the resource pool configuration, you should familiarize yourself with the contents of the configuration file. To display the file, use the `poolcfg` command with the `-c` option, followed by the `info` subcommand, as shown in the example in the slide.

Note

- The `poolcfg` utility is used to create and modify the resource pool configuration files.
- Due to space constraints in the slide, the full output for the resource pool configuration is displayed below.

In the full output, notice that the current pool is `pool_default` and that the processor set (pset) is `pset_default`. Below that, you can see that there are two CPUs associated with the default pset. The number of CPUs available to the pset is identified in the pset value `uint pset.size 2`.

Note: A processor set allows the binding of processes to groups of CPUs.

After checking the resource pool configuration and verifying that it is now the default configuration, you can exit the pool by using the `exit` command.

```

system default
    string    system.comment
    int       system.version 1
    boolean   system.bind-default true
    string    system.poold.objectives wt-load

pool pool_default
    int       pool.sys_id 0
    boolean   pool.active true
    boolean   pool.default true
    int       pool.importance 1
    string    pool.comment
    pset      pset_default

pset pset_default
    int       pset.sys_id -1
    boolean   pset.default true
    uint      pset.min 1
    uint      pset.max 65536
    string    pset.units population
    uint      pset.load 395
    uint      pset.size 2
    string    pset.comment

    cpu
        int      cpu.sys_id 1
        string   cpu.comment
        string   cpu.status on-line

    cpu
        int      cpu.sys_id 0
        string   cpu.comment
        string   cpu.status on-line

```

Modifying the Resource Pool Configuration File

- To create the pset, use `poolcfg -c 'create pset pset_psetname (uint pset.min = x; uint pset.max = x) '`.

```
# poolcfg -c 'create pset pset_1to2 (uint pset.min = 1; uint pset.max = 2) '
```

- To create the pool, use `poolcfg -c 'create pool pool_poolname '`.

```
# poolcfg -c 'create pool pool_hrzone'
```

- To associate the pset with the pool, use `poolcfg -c 'associate pool pool_poolname (pset pset_psetname) '`.

```
# poolcfg -c 'associate pool pool_hrzone (pset pset_1to2) '
```

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

To modify the configuration file to control the CPU usage of `hrzone`, create a pset and then a pool. Associate the pset with the pool.

For example, to address the workload in `hrzone`, two CPUs are allocated to the zone. This enables the kernel to use either one or two CPUs to support the `hrzone` workload.

To create the pset and define its parameters, use the `poolcfg -c` command followed by `create pset`, the name that you want to give the pset (`pset_psetname`), and the unassigned integer (`uint`) minimum and maximum values, which are the minimum and maximum numbers of CPUs that you want allocated to this pset, `pset_1to2`. In the example in the slide, a minimum of one CPU and a maximum of two CPUs are being used.

Note: The `-c` option is used to specify a command.

To create the pool, use the `poolcfg -c` command again, followed by `create pool` and the name that you want to give the pool (`pool_poolname`), as shown in the second example.

To associate the pset with the pool, use the `poolcfg -c` command followed by `'associate pool pool_poolname (pset pset_psetname) '`, as shown in the third example.

Note: Psets and pools are created separately to provide flexibility. For example, If you like, you can create another pset and associate it with `pool_hrzone`.

Displaying and Committing the Modified Resource Pool Configuration File

- To display the modified resource pool configuration, use `poolcfg -c info`.

```
# poolcfg -c info  
<Output presented in the Notes>
```

- To validate and commit the modified configuration, use `pooladm -n -c`, and then use `pooladm -c`.

```
# pooladm -n -c  
# pooladm -c
```



Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

To view the modified pool configuration, use the `poolcfg -c` command followed by the `info` subcommand, as shown in the first example in the slide.

Note: Due to space constraints in the slide, the full output for the resource pool configuration is displayed below.

In the full output, notice that you now have a pool called `pool_hrzone` and a pset called `pset_1to2` as part of your resource pool configuration.

After viewing the modified pool configuration, it is a good practice to validate and commit it. To validate the configuration, use the `pooladm` command with the `-n -c` options, as shown in the second example. After validating, you can commit the configuration by using the `pooladm -c` command. This is your static resource pool configuration file.

```

system default
    string      system.comment
    int         system.version 1
    boolean     system.bind-default true
    string      system.poold.objectives wt-load

pool pool_default
    int         pool.sys_id 0
    boolean     pool.active true
    boolean     pool.default true
    int         pool.importance 1
    string      pool.comment
    pset        pset_default

pool pool_hrzone
    boolean     pool.active true
    boolean     pool.default false
    string      pool.scheduler FSS
    int         pool.importance 1
    string      pool.comment
    pset        pset_1to2

pset pset_default
    int         pset.sys_id -1
    boolean     pset.default true
    uint        pset.min 1
    uint        pset.max 65536
    string      pset.units population
    uint        pset.load 388
    uint        pset.size 2
    string      pset.comment

```

```

cpu
    int      cpu.sys_id 1
    string   cpu.comment
    string   cpu.status on-line

cpu
    int      cpu.sys_id 0
    string   cpu.comment
    string   cpu.status on-line

pset pset_1to2
    int      pset.sys_id -2
    boolean  pset.default false
    uint     pset.min 1
    uint     pset.max 2
    string   pset.units population
    uint     pset.load 0
    uint     pset.size 0
    string   pset.comment

```

Displaying the Resource Pool Configuration That Is Currently in Use

Use `poolcfg -dc info`.

```
# poolcfg -dc info

system default
    string  system.comment
    int     system.version 1
    boolean system.bind-default true
    string  system.poolid.objectives wt-load

pool pool_hrzone
    int     pool.sys_id 1
    boolean pool.active true
    boolean pool.default false
    int     pool.importance 1
    string  pool.comment
    pset    pset_1to2
```

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

To view the resource pool configuration that the system is currently using (that is, the dynamic pool configuration), use the `poolcfg` command with the `-dc` option, followed by the `info` subcommand, as shown in the example in the slide. Observe that `pool_hrzone` is the resource pool configuration that is currently in use.

Note: The `-d` option specifies the dynamic pool configuration (that is, the configuration that is operating directly on the kernel state).

Displaying All Active Resource Pools

To display all the active resource pools on the system, use `poolstat -r all`.

```
# poolstat -r all
```

id	pool	type	rid	rset	min	max	size	used	load
1	pool_hrzone	pset	1	pset_1to2	1	2	1	0.00	0.00
0	pool_default	pset	-1	pset_default	1	66K	1	0.00	0.17

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

The `poolstat` utility is used to report statistics on active resource pools. To display all the active resource pools on the system, use the `poolstat -r all` command, as shown in the example in the slide.

Observe that there are two active resource pools: the pool that was just created (`pool_hrzone`) and the default pool (`pool_default`).

The output for this command displays the following information:

- Pool ID
- Name of the pool
- Type of resource set. A resource set is a process-bindable resource. Examples of resource sets include processor sets and scheduling classes.
- Resource set ID (`rid`)
- Resource set name (`rset`)
- Minimum resource set size (`min`)
- Maximum resource set size (`max`)

- Current resource set size (`size`)
- Amount of the resource set that is currently in use (`used`). This value is calculated as the percentage usage of the resource set multiplied by its size. If the resource set was reconfigured during the last sampling interval, this value might not be reported (`-`).
- Load that is put on the resource set (`load`). For the definition of this property, see `libpool (3LIB)`.

Binding the Zone to a Persistent Resource Pool

1. Allocate the pool to the zone, and then confirm the allocation.
2. Reboot the zone to activate the resource pool binding.
3. Confirm the availability of the resource pool.

The Oracle logo, consisting of the word "ORACLE" in a white, sans-serif font, is centered on a solid red rectangular background.

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

After creating the resource pool, you need to bind the zone to the persistent resource pool. Persistent resource pools remain even if the system shuts down and comes back up.

You can also configure temporary resource pools, which is covered in the section titled “Resource Pools” in *Oracle Solaris Administration: Oracle Solaris Zones, Oracle Solaris 10 Zones, and Resource Management*.

Allocating the Pool to the Zone

- To allocate the pool to the zone, use `zonecfg -z` followed by the zone name, setting `pool=pool_poolname`.

```
# zonecfg -z hrzone set pool=pool_hrzone
```

- To confirm that the allocation has been made, use `zonecfg -z zonename info pool`.

```
# zonecfg -z hrzone info | grep pool
pool: pool_hrzone
```

The Oracle logo, consisting of the word "ORACLE" in a white, sans-serif font, centered on a red rectangular background.

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

After verifying that `hrzone` is running, you can allocate the pool to that zone. To do this, use the `zonecfg -z` command followed by the zone name and `set pool=pool_poolname`, as shown in the first example in the slide.

To confirm that the pool allocation has been included in the zone configuration, use the `zonecfg -z` command followed by the zone name, the `info` subcommand, and `| grep pool`, as shown in the second example.

Rebooting the Zone

- To reboot the zone, use `zoneadm -z hrzone shutdown`.

```
# zoneadm -z hrzone shutdown -r
```

- To verify that the zone is back up and running, use `zoneadm list -iv`.

```
# zoneadm list -iv
```

ID	NAME	STATUS	PATH	BRAND	IP
0	global	running	/	solaris	shared
1	hrzone	running	/zones/hrzone	solaris	excl
2	itzone	running	/zones/itzone	solaris	excl

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

The next step is to reboot the zone to activate the resource pool binding. To do this, use the `zoneadm` command followed by the zone name and the `shutdown` subcommand with the `-r` option, as shown in the first example in the slide.

Note: You can also use `init 6` to reboot.

Then verify that the zone is back up and running. To do this, run the `zoneadm list -iv` command again, as shown in the second example. Observe that `hrzone` is back up and running.

Confirming the Availability of the Resource Pool

Log in by using `zlogin zonename`. Then use `poolcfg -dc info`.

```
# zlogin hrzone
[Connected to zone 'hrzone' pts/2]
Oracle Corporation SunOS 5.11 11.1 September 2012

# poolcfg -dc info
    <Output presented in the Notes>

# exit
Logout
[Connection to zone 'hrzone' pts/2 closed]
```

The Oracle logo, consisting of the word "ORACLE" in a white, sans-serif font, centered on a red rectangular background.

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

To confirm that the resource pool is available, log in to the zone by using the `zlogin` command followed by the zone name. Then use the `poolcfg` command with the `-dc` options, followed by the `info` subcommand, as shown in the example in the slide.

Note: Due to space constraints in the slide, the full output for the resource pool configuration is displayed below.

In the full output, notice a pool called `pool_hrzone` with the pset `pset_1to2`. Below that, you can see the details for the pset as well as the two CPUs that you specified.

After confirming the availability of the resource pool and reviewing the resource pool configuration, you can exit the pool by using the `exit` command.

```

system default
    string      system.comment
    int         system.version 1
    boolean     system.bind-default true
    string      system.poold.objectives wt-load

pool pool_hrzone
    int         pool.sys_id 1
    boolean     pool.active true
    boolean     pool.default false
    int         pool.importance 1
    string      pool.comment
    pset        pset_1to2

pset pset_1to2
    int         pset.sys_id 1
    boolean     pset.default false
    uint        pset.min 1
    uint        pset.max 2
    string      pset.units population
    uint        pset.load 24
    uint        pset.size 1
    string      pset.comment

cpu
    int         cpu.sys_id 0
    string      cpu.comment
    string      cpu.status on-line

```

Removing the Resource Pool Configuration

1. Remove the pool configuration from the zone.
2. Reboot the zone.
3. Check the resource pool configuration for the zone.
4. Delete the resource pool.

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Removing the Resource Pool Configuration from the Zone

Use `zonecfg -z zonename clear pool`.

```
# zonecfg -z hrzone clear pool
```

The Oracle logo, consisting of the word "ORACLE" in a white, sans-serif font, is positioned on a solid red rectangular background.

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

To remove the resource pool configuration from the zone, use the `zonecfg -z` command followed by the zone name and the `clear pool` subcommand, as shown in the example in the slide. This action removes the resource pool configuration that was allocated to the zone and replaces it with the default resource pool configuration.

Rebooting the Zone

- To reboot the zone, use `zoneadm -z hrzone shutdown`.

```
# zoneadm -z hrzone shutdown -r
```

- To verify that the zone is back up and running, use `zoneadm list -iv`.

```
# zoneadm list -iv
```

ID	NAME	STATUS	PATH	BRAND	IP
0	global	running	/	solaris	shared
1	hrzone	running	/zones/hrzone	solaris	excl
2	itzone	running	/zones/itzone	solaris	excl

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

The next step is to reboot the zone to activate the change. To do this, use the `zoneadm` command followed by the zone name and the `shutdown` subcommand with the `-r` option, as shown in the first example in the slide.

You then verify that the zone is back up and running. To do this, run the `zoneadm list -iv` command again, as shown in the second example. Observe that `hrzone` is back up and running.

Checking the Resource Pool Configuration for the Zone

Log in by using `zlogin zonename`. Then use `poolcfg -dc info`.

```
# zlogin hrzone
[Connected to zone 'hrzone' pts/2]
Oracle Corporation SunOS 5.11 11.0 September 2012

# poolcfg -dc info
    <Output presented in the Notes>

# exit
Logout
[Connection to zone 'hrzone' pts/2 closed]
```



Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

To check that the resource pool has been removed, log in to the zone by using the `zlogin` command followed by the zone name. Then use the `poolcfg` command with the `-dc` options, followed by the `info` subcommand, as shown in the example in the slide.

Note: Due to space constraints in the slide, the full output for the resource pool configuration is displayed below.

In the full output, notice that only the default resource pool configuration is available. The only pool that is available is the `pool_default` pool, and the only pset that is available is `pset_default`. Below the default pset, you can see the two CPUs that are associated with that pset.

After checking the resource pool configuration and verifying that it is now the default configuration, you can exit the pool by using the `exit` command.

```
system default
    string      system.comment
    int         system.version 1
    boolean     system.bind-default true
    string      system.poold.objectives wt-load

pool pool_default
    int         pool.sys_id 0
    boolean     pool.active true
    boolean     pool.default true
    int         pool.importance 1
    string      pool.comment
    pset        pset_default

pset pset_default
    int         pset.sys_id -1
    boolean     pset.default true
    uint        pset.min 1
    uint        pset.max 65536
    string      pset.units population
    uint        pset.load 268
    uint        pset.size 1
    string      pset.comment

cpu
    int         cpu.sys_id 1
    string      cpu.comment
    string      cpu.status on-line
```


Deleting the Resource Pool

- To delete the resource pool, use `pooladm -x`.

```
# pooladm -x
```

- To display all the active resource pools on the system, use `poolstat -r all`.

```
# poolstat -r all
id pool          type rid rset          min  max   size   used   load
0 pool_default  pset  -1 pset_default  1    66K   2     0.00  0.73
```

The Oracle logo, consisting of the word "ORACLE" in a white, sans-serif font on a red rectangular background.

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

The next step is to delete the resource pool. To do this, use the `pooladm` command with the `-x` option, as shown in the first example in the slide. The `-x` option removes the currently active pool configuration, destroys all defined resources, and returns all formerly partitioned components to their default resources.

To verify that the deleted resource pool is no longer active on the system, use the `poolstat -r all` command, as shown in the second example in the slide. Observe that there is now only one active resource pool: the default pool (`pool_default`).

Resource Capping

- A memory resource conflict in your system can be addressed by capping the amount of memory that is allocated to each zone.
- Resource capping is controlled by the `rcapd` daemon.
- The `rcapd` daemon repeatedly samples the resource utilization of projects that have physical memory caps.
- The sampling interval is specified by the administrator.
- When physical memory utilization thresholds are breached, the daemon reduces the resource consumption with memory caps.

The Oracle logo, consisting of the word "ORACLE" in a white, sans-serif font, is centered within a solid red rectangular bar.

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Note: You can use the `rcapadm` command without arguments to display the current status of the resource capping daemon.

For more information about resource capping and the `rcapd` daemon, see the section titled “Administering the Resource Capping Daemon” in *Oracle Solaris Administration: Oracle Solaris Zones, Oracle Solaris 10 Zones, and Resource Management*.

Allocating Physical Memory Resources with Resource Capping

To add a memory cap to a zone:

1. Configure the zone by using `zonecfg -z zone`.
2. Add the `capped-memory` resource, and set each memory cap type: `physical`, `swap`, and `locked`.
3. Verify, commit, and exit the zone.

```
# zonecfg -z hrzone
zonecfg:hrzone> add capped-memory
zonecfg:hrzone:capped-memory> set physical=50m
zonecfg:hrzone:capped-memory> set swap=100m
zonecfg:hrzone:capped-memory> set locked=30m
zonecfg:hrzone:capped-memory> end
zonecfg:hrzone> verify
zonecfg:hrzone> commit
zonecfg:hrzone> exit
```

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

To allocate the maximum amount of memory that can be consumed by a specified zone and have it as a persistent cap, you can use the `capped-memory` resource. The `capped-memory` resource sets limits for `physical`, `swap`, and `locked` memory. Each limit is optional, but at least one limit must be set.

Note: Ensure that the `rcapd` daemon and `rcap` service are up and running.

In the example, the `physical`, `swap`, and `locked` memory resource caps have been set for `hrzone`.

Quiz

You do not necessarily need to reboot the zone to activate the resource pool binding. Validating and committing the modified resource pool configuration file by using `pooladm -n -c` can activate the resource pool binding.

- a. True
- b. False

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Answer: b

Lesson Agenda

- Explaining the fundamentals of Oracle Solaris Zones
- Configuring Oracle Solaris Zones
- Configuring network connectivity in Oracle Solaris Zones
- Administering Oracle Solaris Zones
- Managing system resources in Oracle Solaris Zones
- **Securing Oracle Solaris Zones**

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Securing Oracle Solaris Zones

- By design, Oracle Solaris Zones are secure from external attacks and internal malicious programs.
- Features such as the following provide some inherent restrictions and security mechanisms for zones:
 - Delegated administration
 - Zone link protection
 - Exclusive IP
 - Immutable zones (read-only root)
 - Cryptographic services
 - Privileges
 - Users and rights profiles

The Oracle logo, consisting of the word "ORACLE" in a white, sans-serif font, is centered within a solid red rectangular bar.

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

You learn more about privileges and rights profiles in the lesson titled “Administering User Accounts.”

For more information about Oracle Solaris Zones security, see the *Oracle Solaris 11 Security Administration* course.

Delegated Administration

- Oracle Solaris enables the delegation of common administrative tasks for specific zones to specific administrators by using RBAC.
- This example shows user `john` being given the authorization to manage a specific zone called `hrzone`:

```
# zonecfg -z hrzone
zonecfg:zoneA> add admin
zonecfg:zoneA> set user=john
zonecfg:zoneA> set auths=manage
zonecfg:zoneA> end
zonecfg:zoneA> commit
zonecfg:zoneA> exit
#
```



Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

With Oracle Solaris 11, you have the ability to delegate common zone administration tasks for specific zones to different administrators. For each zone, a user (or set of users) can be identified with the permissions to log in, manage, or clone from that zone.

These specific authorizations are interpreted by the appropriate commands running in the global zone to allow access at the correct authorization level to the correct user. There are two basic levels of zone administration:

- **Global administrator:** A global administrator has superuser privileges or an equivalent rights profile. When logged in to the global zone, the global administrator can monitor and control the system as a whole.
- **Zone administrator:** A non-global zone can be administered by a zone administrator. The global administrator assigns the required authorizations to the zone administrator. The privileges of a zone administrator are confined to a non-global zone.

Being able to delegate administration is extremely useful in a shared environment where you want to allow specific users to be able to manage the zones that are relevant to their role.

Zone Link Protection

- Oracle Solaris 11 provides a new link protection mechanism to prevent potentially malicious or misbehaving zones from sending harmful packets to the network.
- Link protection shields against basic threats, including IP, DHCP, MAC, and L2 frame spoofing.

The Oracle logo, consisting of the word "ORACLE" in a white, sans-serif font, is centered on a solid red rectangular background.

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

In virtualized environments, it is common for the host administrator to grant to a zone exclusive access to a physical link or a VNIC. Doing so enables zones to benefit from traffic isolation and improved performance.

However, the zones can generate any type of traffic (even harmful packets) and send that traffic over the network. The new zone link protection mechanism in Solaris 11 helps prevent the threat of harmful traffic.

Exclusive IP

- In virtualized environments, a dedicated physical network interface controller is needed to achieve network separation between virtual environments.
- Exclusive IP is the default networking configuration for non-global zones.
- Exclusive IP in Oracle Solaris gives 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.

The Oracle logo, consisting of the word "ORACLE" in a white, sans-serif font, centered within a red rectangular background.

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

In a shared-IP zone, processes cannot send packets with source IP addresses other than the ones that are assigned by using the `zonecfg` utility.

A shared-IP zone does not have access to send and receive arbitrary datalink (layer 2) packets. For an exclusive-IP zone, `zonecfg` grants the entire specified datalink to the zone. Someone who compromises an exclusive-IP zone can forge IP addresses (which can also happen in the global zone).

In addition, you can run IPsec and IKE in an exclusive-IP zone but not in a shared-IP zone. IPsec policy can include IP addresses for shared-IP zones. IPsec can be configured to support shared-IP zones, but all configuration must be performed in the global zone. You cannot use IKE to manage keys in a shared-IP zone.

You can further secure zones through IP Filter. Oracle Solaris IP Filter, including Network Address Translation (NAT) functionality, can be run in an exclusive-IP zone. You cannot run IP Filter from within a shared-IP zone. IP Filter cannot perform filtering on packets between zones on the same system. When using IP Filter to filter between zones, you must add the following line to the top of an IP Filter rule set:

```
set intercept_loopback true;
```

Immutable Zones

- Oracle Solaris supports the creation of *immutable* zones (zones that have read-only file systems).
- The `file-mac-profile` property is used to configure a read-only zone root.
- The `zonecfg` utility is used to set the `file-mac-profile` property, as in the following examples:
 - Setting a strict read-only zone

```
zonecfg:hrzone: set file-mac-profile=strict
```

- Setting a fixed-configuration read-only zone

```
zonecfg:hrzone: set file-mac-profile=fixed-configuration
```

- Setting a flexible-configuration read-only zone

```
zonecfg:hrzone: set file-mac-profile=flexible-configuration
```

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

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 that are running in immutable zones.

The `file-mac-profile` property specifies which part of the file system is exempt from the read-only policy. There are four possible values:

- **none**: Makes the zone exactly the same as a normal read/write zone. Setting the value to `none` is equivalent to not setting the `file-mac-profile` property.
- **strict**: Allows no exceptions to the read-only policy
- **fixed-configuration**: Allows the zone to write to files in and below `/var` (except directories containing configuration files)
- **flexible-configuration**: Similar to `fixed-configuration` but also allows writing to files in `/etc`

Cryptographic Services in Zones

- The global zone and each non-global zone have their own `/system/cryptosvc` services.
- When the cryptographic service is enabled or refreshed in the global zone:
 - The `kcfld` daemon starts in the global zone
 - The user-level policy for the global zone is set
 - The kernel policy for the system is set
- When the service is enabled or refreshed in a non-global zone:
 - The `kcfld` daemon starts in the zone
 - The user-level policy for the zone is set
 - The kernel policy is set by the global zone

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Privileges

- The implementation of *least privilege* in Oracle Solaris provides a set of fine-grained privileges to replace the concept of the all-powerful `root` role.
- When it is booted, a zone provides a default set of safe privileges.
- Processes are restricted to a subset of privileges.
- Privilege restriction prevents a zone from performing operations that might affect other zones.
- *Privilege sets* limit the capabilities of privileged users within the zone.
- To display the list of privileges that are available from within a given zone, use the `ppriv` utility.

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Instead of performing checks against `UID 0` to allow privileged operations, the kernel now checks for the specific privileges that are required to perform privileged operations.

In the past, it was sufficient to be the superuser to perform mount operations. Now, even the `root` role must have a specific privilege to perform mount operations.

By restricting the privileges of `root` in the non-global zone to a set of privileges that are safe in a zone, the `root` role in a non-global zone can be given enough privileges to manage the zone without the ability to affect the system as a whole. For example, a non-global zone user must not be able to reboot the system.

Restricting privileges is not sufficient for isolation. Privileges restrict only the operations that can be performed, and not the objects on which they are performed. This is accomplished by the isolation that zones provide.

It is therefore possible to delegate non-global zone administration to users by giving them access to the `root` account in a non-global zone. Because a user with `uid 0` in one zone is different from a user with `uid 0` in another zone, a non-global zone administrator cannot compromise any other zone. However, the global zone `root` role must be closely monitored, because this user has control over the system as a whole and thus has access to all zones.

Privileges

- A zone's privilege set specifies the upper bound of privileges that the zone and its processes can obtain.
- There are four groups of privileges in Oracle Solaris Zones:
 - **Default privileges** are available, by default, in a zone's privilege set.
 - **Required privileges** must be listed in a zone's privilege set.
 - **Prohibited privileges** cannot be included in the privilege set of a zone.
 - **Optional privileges** are not part of the default set of privileges granted to a zone, but they can be added to a zone's privilege set if required.

The Oracle logo, consisting of the word "ORACLE" in a white, sans-serif font, is centered on a solid red rectangular background.

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

The default set is smaller than the set of privileges in the global zone. The status of a privilege in a zone (default, required, prohibited, or optional) dictates whether the privilege can be assigned through the `limitpriv` property of `zonecfg(1M)`.

Use the following command to display the privileges of the `root` role in a non-global zone:

```
# ppriv -v $$
```

For information about privileges in zones, see the `privileges(5)` man page.

Users and Rights Profiles in Oracle Solaris Zones

- Rights profiles cannot grant permission for actions (such as setting the system time) that are not allowed in zones.
- If a rights profile allows access to a global zone-only process, the user or role with that rights profile will still be unable to perform the action.
- If a privilege is granted in the `exec_attr` or `prof_attr` databases but is not part of the zone's privilege set, the privilege cannot be granted and will result in a failure.

Note: To administer zones, you must be assigned Zone Management or Zone Security rights profiles.

The Oracle logo, consisting of the word "ORACLE" in a white, sans-serif font, is positioned on a solid red rectangular background.

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Summary

In this lesson, you should have learned how to:

- Explain the fundamentals of Oracle Solaris Zones
- Configure Oracle Solaris Zones
- Configure network connectivity in Oracle Solaris Zones
- Administer Oracle Solaris Zones
- Manage system resources in Oracle Solaris Zones
- Secure Oracle Solaris Zones


ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Administering Privileges and RBAC

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.



System Administration for Experienced UNIX/Linux Administrators



**Administering System
Software by Using IPS**



**Administering Services
by Using SMF**



Administering ZFS



Configuring the Network



**Administering Oracle Solaris
Zones**



**Administering Privileges
and RBAC**



**Installing the Oracle Solaris 11
Operating System**



Monitoring System Resources

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Objectives

After completing this lesson, you should be able to:

- Administer process rights management
- Configure Role-Based Access Control (RBAC)

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Lesson Agenda

- Administering process rights management
- Configuring RBAC

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Assignment of User Privileges and Roles

- As an administrator, you must ensure that processes and users have only the level of access or privilege to system resources that are needed to perform required functions—and no more.
- To do this, address the following specific areas:
 - Process rights management
 - Role-Based Access Control (RBAC)

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Process Rights Management

- Process rights management
 - Enables processes to be restricted at the command, user, role, or system levels
 - Is implemented through privileges
- Privileges
 - Decrease the security risk that can occur if one user or one process has full superuser capabilities on a system
 - Allow a gradation between user capabilities and root capabilities
 - Restrict programs and processes to only those capabilities that the program requires (this is the principle of “least privilege”)

The Oracle logo, consisting of the word "ORACLE" in a white, sans-serif font, is centered within a solid red rectangular bar.

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

A system that enforces policy with privileges allows a gradation between user capabilities and root capabilities. A user can be granted privileges to perform activities that are beyond the capabilities of regular users, and `root` can be limited to fewer privileges than `root` currently possesses. With RBAC, a command that runs with privileges can be isolated in a rights profile and assigned to a user or role.

Least Privilege

On a system that implements least privilege, an intruder who captures a process can access only those privileges that the process has. The rest of the system is not compromised.

Areas of Privilege

- **FILE** privileges
- **IPC** privileges
- **NET** privileges
- **PROC** privileges
- **SYS** privileges

The Oracle logo, consisting of the word "ORACLE" in a white, sans-serif font, centered within a red rectangular background.

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Privileges are logically grouped on the basis of the area of the privilege.

- **FILE privileges:** Privileges that begin with the string `file` operate on file system objects. For example, the `file_dac_write` privilege overrides discretionary access control when writing to files.
- **IPC privileges:** Privileges that begin with the string `ipc` override IPC object access controls. For example, the `ipc_dac_read` privilege enables a process to read remote shared memory that is protected by DAC.
- **NET privileges:** Privileges that begin with the string `net` give access to specific network functionality. For example, the `net_rawaccess` privilege enables a device to connect to the network.
- **PROC privileges:** Privileges that begin with the string `proc` enable processes to modify restricted properties of the process itself. PROC privileges include privileges that have a very limited effect. For example, the `proc_clock_highres` privilege enables a process to use high-resolution timers.
- **SYS privileges:** Privileges that begin with the string `sys` give processes unrestricted access to various system properties. For example, the `sys_linkdir` privilege enables a process to make and break hard links to directories.

Sets of Privileges

Every process has the following four sets of privileges that determine whether the process can use a particular privilege:

- Effective privilege set (\mathbb{E})
- Inheritable privilege set (\mathbb{I})
- Permitted privilege set (\mathbb{P})
- Limit privilege set (\mathbb{L})

The Oracle logo, consisting of the word "ORACLE" in a white, sans-serif font, is centered on a solid red rectangular background.

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

The kernel automatically calculates the effective set of privileges. You can modify the initial inheritable set of privileges. A program that is coded to use privileges can reduce the program's permitted set of privileges. You can shrink the limit set of privileges.

- **Effective privilege set (\mathbb{E}):** The set of privileges that are currently in effect. A process can add privileges that are in the permitted set to the effective set. A process can also remove privileges from \mathbb{E} .
- **Inheritable privilege set (\mathbb{I}):** The set of privileges that a process can inherit across a call to `exec`. After the call to `exec`, the permitted and effective sets are equal, except in the special case of a `setuid` program. For a `setuid` program, after the call to `exec`, the inheritable set is first restricted by the limit set. Then the set of privileges that were inherited (\mathbb{I}), minus any privileges that were in the limit set (\mathbb{L}), are assigned to \mathbb{P} and \mathbb{E} for that process.

- **Permitted privilege set (P):** The set of privileges that are available for use. Privileges can be available to a program from inheritance or through assignment. An execution profile is one way to assign privileges to a program. The `setuid` command assigns all privileges that `root` has to a program. Privileges can be removed from the permitted set, but privileges cannot be added to the set. Privileges that are removed from `P` are automatically removed from `E`. **Note:** A privilege-aware program removes the privileges that a program never uses from the program's permitted set. In this way, unnecessary privileges cannot be exploited by the program or by a malicious process.
- **Limit privilege set (L):** The outside limit of privileges that are available to a process and its children. By default, the limit set is all privileges. Processes can shrink the limit set but can never extend the limit set. `L` is used to restrict `I`. As a result, `L` restricts `P` and `E` at the time of execution.

If a user is assigned a profile with a program that has been assigned privileges, the user can usually run that program. On an unmodified system, the program's assigned privileges are within the user's limit set. The privileges that have been assigned to the program become part of the user's permitted set. To run the program that has been assigned privileges, the user must run the program from a profile shell.

The kernel recognizes a basic privilege set. On an unmodified system, each user's initial inheritable set equals the basic set at login. Although you cannot modify the basic set, you can modify the privileges that a user inherits from the basic set. At login, all users have the basic set in their inheritable set, their permitted set, and their effective set. A user's limit set is equivalent to the default limit set for the zone (global or non-global). To put more privileges in the user's effective set, you must assign a rights profile to the user. The rights profile would include commands to which you have added privileges.

Administering Privileges

Administering privileges involves managing the following:

- Process privileges
 - Determining the privileges that are available to the shell
 - Determining the privileges on a process
 - Displaying the description of a privilege
- User privileges
 - Determining the privileges that are directly assigned to you
 - Determining the privileged commands that a user can use
 - Assigning privileges to a user or role
 - Limiting the privileges of a user or role
 - Debugging privilege failure

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

User Privileges

The most secure way to manage privileges for users and roles is to confine the use of a privilege to commands in a rights profile. The rights profile is then included in a role, and the role is assigned to a user. When the user assumes the assigned role, the privileged commands are available to be run in a profile shell.

Determining the Privileges Available to the Shell

To determine the privileges that are available to your processes, use `ppriv $$` to list the process privileges that are available to your shell.

```
# ps
  PID TTY          TIME CMD
  990 pts/1        0:00 bash
  993 pts/1        0:00 ps
# ppriv $$
990:  bash
flags = <none>
      E: all
      I: basic
      P: all
      L: all
```

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

The `ppriv` command is used to inspect or modify process privilege sets and attributes. The double-dollar sign (`$$`) passes the process number of the parent shell to the command.

In the example, the `ps` command displays the processes that are currently running and the shell used. Note that the shell in use is the `bash` shell. The `ppriv $$` command also indicates that the shell is `bash`. There are no flags set; the effective (`E`), permitted (`P`), and limit (`L`) privilege sets are set to `all`; and the inherited (`I`) privilege set is set to `basic`.

Note: The `flags` field is associated with the `getpflags()` and `setpflags()` functions that are used to get or set process flags. The following values are supported:

- **PRIV_AWARE:** This one-bit flag takes the value of 0 (unset) or 1 (set). The current process becomes privilege-aware only when this flag is set. See `privileges(5)` for a discussion of this flag.
- **PRIV_DEBUG:** This one-bit flag takes the value of 0 (unset) or 1 (set). The current process has privilege debugging enabled only when this flag is set.
- **NET_MAC_AWARE** and **NET_MAC_AWARE_INHERIT:** These flags are available only if the system is configured with Trusted Extensions (TX). These one-bit flags each take the value of 0 (unset) or 1 (set).

Determining the Privileges Available to the Shell

To display the names of the privileges in each privilege set, use `ppriv -v $$`.

```
# ppriv -v $$
990:~$
flags = <none>
E: contract_event,contract_identity,contract_observer,cpc_cpu,dtrace_kernel,
dtrace_proc,dtrace_user,file_chown,file_chown_self,file_dac_execute,
<output omitted>
I: file_link_any,file_read,file_write,net_access,proc_exec,proc_fork,
proc_info, proc_session
P: contract_event,contract_identity,contract_observer,cpc_cpu,dtrace_kernel,
dtrace_proc,dtrace_user,file_chown,file_chown_self,file_dac_execute,
<output omitted>
L: contract_event,contract_identity,contract_observer,cpc_cpu,dtrace_kernel,
dtrace_proc,dtrace_user,file_chown,file_chown_self,file_dac_execute,
<output omitted>
```

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

The example in the slide illustrates using the `-v` option with the `ppriv $$` command to display the names of privileges by privilege set (this example shows only partial output). Take a closer look at the privileges in the inheritable (I) privilege set. The privileges listed there indicate that you can link to any file, read any file, and write any file. You have access to the network, which means you can perform network configuration tasks. In addition, you can execute any process, run a process in another subshell (`proc_fork`), display information about any processes, and view any session in the process.

Determining the Privileges on a Process

To determine the privileges that are available to a process, use `ppriv -v pid`.

```
# ppriv -v 476
476: /usr/sbin/cron
flags = <none>
E: contract_event,contract_identity,contract_observer,cpc_cpu,
   dtrace_kernel,dtrace_proc,dtrace_user,file_chown,
<output omitted>
I: file_link_any,file_read,file_write,net_access,proc_exec,
   proc_fork,proc_info,proc_session
P: contract_event,contract_identity,contract_observer,
   cpc_cpu,dtrace_kernel,dtrace_proc,dtrace_user,file_chown,
<output omitted>
L: contract_event,contract_identity,contract_observer,
   cpc_cpu,dtrace_kernel,dtrace_proc,dtrace_user,file_chown,
<output omitted>
```



Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Use the `ppriv -v` command with the process ID number (PID). The example shows partial output for the `cron` process.

Displaying the Description of a Privilege

To display a privilege definition, use `ppriv -vl <privilege>`.

```
# ppriv -vl contract_event
contract_event
    Allows a process to request critical events without
    limitation.
    Allows a process to request reliable delivery of all
    events on any event queue.
# ppriv -vl proc_exec
proc_exec
    Allows a process to call execve().
#
```

The Oracle logo, consisting of the word "ORACLE" in a white, sans-serif font on a red rectangular background.

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

If you need the definition of a privilege that is listed for a process, use the `ppriv -vl` command followed by the privilege name.

The first example is for the `contract_event` privilege, and the second is for the `proc_exec` privilege.

Determining the Privileges That Are Directly Assigned to You

To view the privileges that have been directly assigned to your user account, use `ppriv -v $$`.

```
# ppriv -v $$
990:  bash
flags = <none>
  E: file_link_any,proc_clock_highres,proc_session
  I: file_link_any,proc_clock_highres,proc_session
  P: file_link_any,proc_clock_highres,proc_session
  L: cpc_cpu,dtrace_kernel,dtrace_proc,dtrace_user,sys_time
# ppriv -vl proc_clock_highres
  Allows a process to use high resolution timers.
```

The Oracle logo, consisting of the word "ORACLE" in a white, sans-serif font, is centered on a solid red rectangular background.

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

The privileges that are listed in the effective set are in effect throughout your session. If you have been directly assigned privileges in addition to the basic set, the privileges are listed in the effective set.

In this example, the user always has access to the `proc_clock_highres` privilege. This privilege allows a process to use high-resolution timers.

Note: To see the privileges that have been directly assigned to a role, you `su` to the role and then run the `ppriv -v $$` command.

Determining the Privileged Commands That a User Can Use

To determine the rights profiles that have been assigned to you, use `profiles`.

```
# profiles
    Basic Solaris User
    All
# profiles -l
    All
    *
    Basic Solaris User
    /usr/bin/cdda2wav.bin
    privs=file_dac_read,sys_devices,proc_prioctl,net_privaddr
    /usr/bin/cdrecord.bin
    privs=file_dac_read,sys_devices,proc_lock_memory,proc_prioctl,net_privaddr
    /usr/bin/readcd.bin    privs=file_dac_read,sys_devices,net_privaddr
```

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

When a user is not directly assigned privileges, the user obtains access to privileged commands through a rights profile. Commands in a rights profile must be executed in a profile shell. Use the `profiles` command to determine the privilege commands that you can use (these are actually the rights profiles that have been assigned to you). To see more details about the privileges, use the `profiles -l` command.

Note: To display the details of a specific privilege, use the `profiles -l` command with the privilege name, as in the following example:

```
# profiles -l Basic Solaris User
```

To display the roles and authorization privileges that you have, use the `roles` and `auth` commands, respectively, as in this example:

```
# roles
```

```
No roles
```

```
# auths
```

```
solaris.admin.wusb.read,solaris.device.cdrw,solaris.device.mount
.removable,solaris.mail.mailq,solaris.profmgr.read
```


Assigning Privileges to a User or Role

- To assign privileges to a user, use `usermod -K key=value <loginname>`.

```
# usermod -K defaultpriv=basic,proc_clock_highres jjones
# getent user_attr | grep jjones
jjones:::type=normal;defaultpriv=basic,proc_clock_highres
```

- To assign privileges to a role, use `rolemod -K key=value <rolename>`.

```
# rolemod -K defaultpriv=basic,proc_clock_highres realtime
# getent user_attr | grep realtime
realtime:::type=role;defaultpriv=proc_clock_highres
```

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

You might want to assign a particular privilege to a user or role all the time. Very specific privileges that affect a small part of the system are good candidates for assigning to a user or role. To assign privileges to a user, use the `usermod -K` command followed by the `key=value` pair that you want to assign and the user's login name.

Note: The `-K key=value` option is used to replace or add to a user's or role's `key=value` pair attributes. See `user_attr (4)` for a list of valid `key=value` pairs.

In the first example in the slide, user `jjones` is enabled to use high-resolution timers by assigning the `proc_clock_highres` privilege to the user's basic default privileges. The values for the `defaultpriv` keyword replace the existing values. Therefore, for `jjones` to retain the `basic` privileges, the value `basic` must be specified. In the default configuration, all users have basic privileges. To verify that the privilege has been assigned, look at the `user_attr` entry for `jjones`. You can see how the privileges have been modified.

To assign privileges to a role, the same method applies. Use the `rolemod -K` command followed by `key=value` pair that you want to assign and the role name. In the second example in the slide, change `user` to `role` as appropriate. The role name in the example is `realtime`.

Limiting the Privileges of a User or Role

1. Determine the privileges in a user's or role's basic set and limit set.
2. Remove one of the privileges from the basic set or from the limit set.
3. Verify that the user or role can still perform other assigned functions as required.

```
# usermod -K limitpriv=all,!sys_linkdir jjones
# getent user_attr | grep jjones
jjones::::type=normal;defaultpriv=basic;limitpriv=all,!sys_linkdir
```

```
# rolemod -K limitpriv=all,!sys_linkdir realtime
# getent user_attr | grep realtime
realtime::::type=role;defaultpriv=basic;limitpriv=all,!sys_linkdir
```



Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

There might be situations when you want to limit the privileges that are available to a user or role. You can do this by reducing the basic set or by reducing the limit set. However, you should have a good reason to limit the privileges, because such limitations can have unintended consequences. To limit the privileges of a user or role, follow the steps listed in the slide.

Caution for step 2: Do not remove the `proc_fork` privilege or the `proc_exec` privilege. Without these privileges, the user cannot use the system. In fact, these two privileges should be removed only from daemons that do not `fork()` or `exec()` other processes.

Notes for step 3: You must thoroughly test a user's or role's capabilities if you have modified its basic set or the limit set. It is possible to prevent a user or role from using the system when the basic set is less than the default. When you modify the limit set to be less than all privileges, processes that require an effective `UID=0` to run might possibly fail.

In the first example, all sessions that originate from `jjones`'s initial login are prevented from using the `sys_linkdir` privilege. After this change is implemented, the user `jjones` can no longer make hard links to directories or unlink directories even after running the `su` command. In the second example, the same consequences apply to the `realtime` role.

Debugging Privilege Failure

Oracle Solaris provides two tools to debug privilege failure:

- `ppriv` command
- `truss` command

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Debugging Privilege Use in a Profile Shell

The `ppriv` command can debug privilege use in a profile shell.

```
jjones:~$ ls -l useful.script
-rw-r--r-- 1 aloa staff 2303 Dec 15 10:10 useful.script
jjones:~$ chown objadmin useful.script
chown: useful.script: Not owner
jjones:~$ ppriv -eD chown objadmin useful.script
chown[11444]: missing privilege "file_chown"
          (euid = 130, syscall = 16) needed at zfs_zaccess+0x258
chown: useful.script: Not owner
```

Note: If you assign a rights profile to a user, and if the rights profile includes commands with privileges, the commands must be entered in a profile shell. If the privileged commands are in a regular shell, the commands do not execute with privilege.

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Observe that `jjones`'s attempt to change the permissions on the `useful.script` file fails. The user then runs the `ppriv` debugging command to determine why the command failed. The output indicates that the `file_chown` privilege is missing. To fix this issue, the `file_chown` privilege needs to be assigned to the `jjones` user.

Quiz

Which two privileges would you remove so that a user cannot use the system?

- a. `proc_fork`
- b. `sys_linkdir`
- c. `proc_exec`
- d. `proc_priocntl`

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Answer: a, c

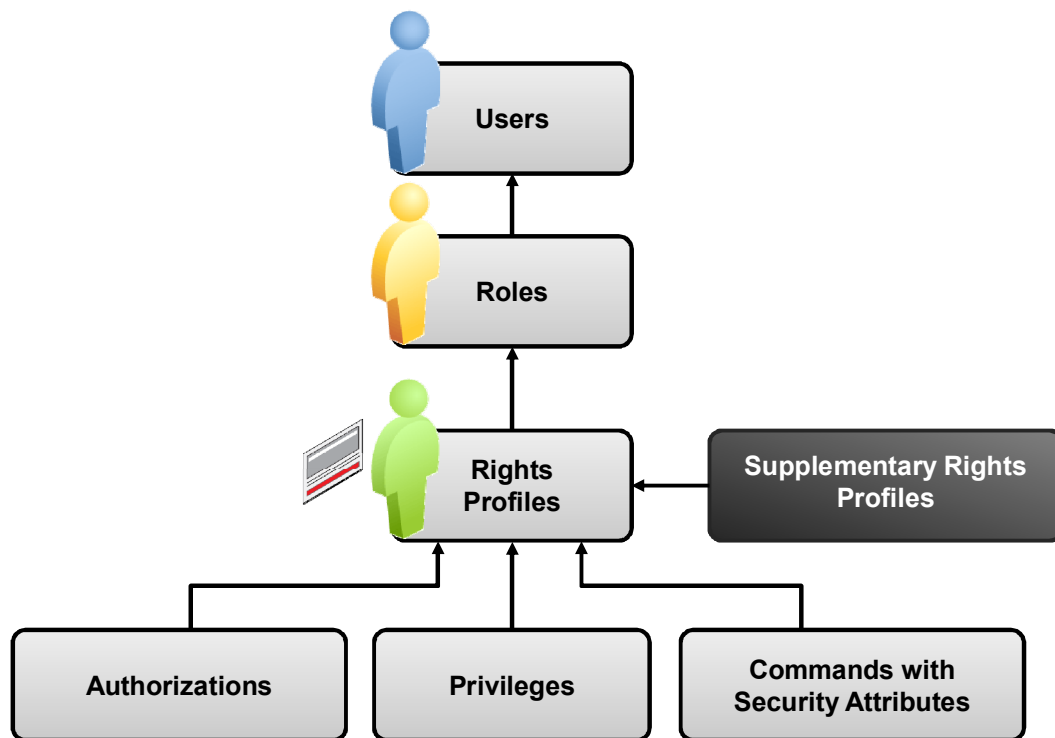
Lesson Agenda

- Administering process rights management
- Configuring RBAC

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Role-Based Access Control (RBAC)

**ORACLE**

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

RBAC is a security feature for controlling user access to tasks that are normally restricted to the superuser. RBAC collects superuser capabilities into rights profiles. Rights profiles can contain:

- Authorizations
- Privileges
- Privileged commands
- Other supplementary rights profiles

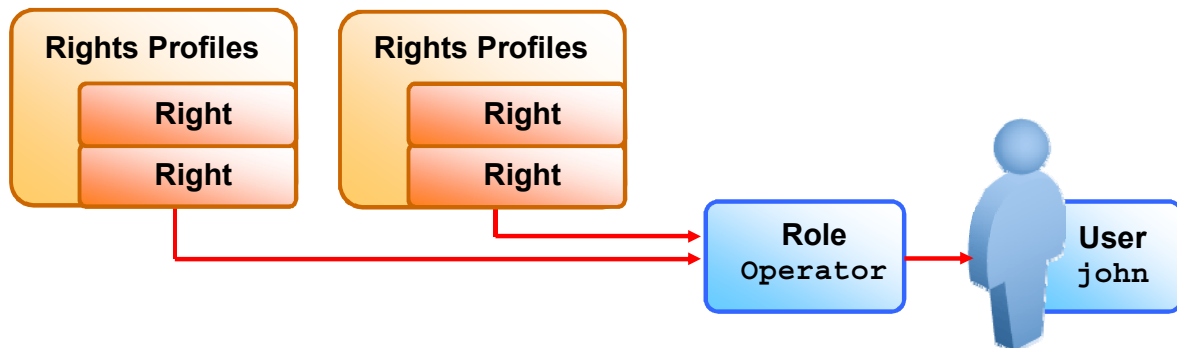
Privileged commands are commands that execute with security attributes.

Rights profiles are assigned to special user accounts that are called *roles*. A user can then assume a role to do a job that requires some of the superuser capabilities.

Roles

A *role*:

- Is a special type of user account that performs a set of administrative tasks
- Contains one or more rights profiles
- Provides access to restricted functionality



```
# roles john
Operator
```

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

A role can be shared among users. This is why roles are preferred in RBAC—because they simplify the management of large numbers of users. Note that a role cannot log in to the system; a user must be logged in to the system to assume a role.

The graphic illustrates user `john` assigned the `Operator` role, which contains several rights profiles. The roles assigned to a user can be displayed by using the `roles` command, as shown in the code example. The user `john` has one assigned role: `Operator`.

Rights Profile

- A *rights profile* is a collection of authorizations, commands, and other rights profiles that are assigned to a user or role.
- *Rights* are commands or scripts that run with special security attributes.
- The profiles that are assigned to a user can be displayed by using the `profiles` command.

```
# profiles john
Operator
Basic Solaris User
All
```

Note: All users by default have the Basic Solaris User profile, which grants users access to all listed authorizations.

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Authorizations

- Authorizations:
 - Enforce policy at the user application level
 - Can be assigned directly to a role or to a user
 - Are typically included in a rights profile
- The rights profile is included in a role, and the role is assigned to a user.
- The authorizations that are assigned to a user can be displayed by using the `auths` command.
 - For example, security policy at installation gives regular users the `solaris.device.cdrw` authorization, which enables them to read and write to a CD-ROM device.

```
# auths john
solaris.admin.wusb.read,solaris.device.cdrw,solaris.device.moun
t.removable,solaris.mail.mailq,solaris.profmgr.read
```

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

An *authorization* is a name associated with the right to access restricted functionality. Note that authorizations can be assigned to user accounts or roles, or they can be embedded in a rights profile, which can then be assigned to a user or a role.

An authorization has a *hierarchy* in which the levels are separated by periods. For example, the `solaris.network.autoconf.read` authorization has the following hierarchy:

- **First level:** `solaris`
- **Second level:** `network.autoconf`
- **Third level:** `read`

This entry gives the basic user the authority to display the rights profile. Similarly, the `solaris.mail.mailq` authorization enables the basic user to look at the mail queue. The default authorizations for every user can be defined in the `/etc/security/policy.conf` file.

In the code example, the authorizations assigned to user `john` are displayed. `john` has all Oracle Solaris authorizations assigned to him.

Privileges

- A *privilege* is a discrete right that can be granted to a command, a user, a role, or a system.
- Privileges enable a process to succeed.
 - **Example:** The `proc_exec` privilege allows a process to call `execve()`.
- Regular users have basic privileges.

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Security Attributes

A security attribute:

- Enables a process to perform an operation that is otherwise forbidden to regular users
- Includes authorizations and privileges
- Can be assigned to a user

ORACLE


Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

In the RBAC model, authorizations and privileges are security attributes (in addition to the `setuid` and `setgid` programs). These attributes can be assigned to a user. For example, a user with the `solaris.device.allocate` authorization can allocate a device for exclusive use.

Privileges can be placed on a process. For example, a process with the `file_flag_set` privilege can set `immutable`, `no-unlink`, or `append-only` file attributes.

Key RBAC Files

The roles, rights profiles, authorizations, privileges, and commands are defined in the following four RBAC files:



ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

- **user_attr:** Contains the rights profiles and authorizations associated with users and roles that supplement the `/etc/passwd` and `/etc/shadow` files
- **auth_attr:** Contains authorization attributes
- **exec_attr:** Contains execution attributes
- **prof_attr:** Contains rights profiles

The relationships among these files are illustrated in the graphic.

Key RBAC Files: `user_attr`

- The `user_attr` file uses colons (:) to separate the fields on each line.
- The first field is the username as it appears in the `/etc/passwd` and `/etc/shadow` files.
- The middle fields are reserved for future use, and the last field is a list of semicolon-separated (;) key-value pairs that describe the security attributes to be applied when the user runs commands.

```
# getent user_attr | grep chris  
chris::::type=normal;profiles=Printer Management
```

The Oracle logo, consisting of the word "ORACLE" in a white, sans-serif font, is positioned on the right side of a solid red horizontal bar.

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Key RBAC Files: auth_attr

```
# getent auth_attr | more
solaris.smf.manage.cups:::Manage CUPS service
states::help=ManageCUPS.html
solaris.smf.manage.dt:::Manage Desktop Service
States::help=ManageDtHeader.html
solaris.smf.manage.dt.login:::Manage Desktop Login Service
States::help=ManageDt
Login.html
solaris.smf.manage.dbus:::Manage D-BUS Service
States::help=SmfDBUSStates.html
solaris.smf.value.tcsd:::Change TPM Administration value properties::
solaris.smf.manage.tcsd:::Manage TPM Administration service states::
solaris.smf.manage.servicetags:::Manage Service Tags Service
States::help=StStat
es.html
solaris.smf.value.servicetags:::Change Service Tag Service Property
Values::help
=StValue.html
solaris.:RO::All Solaris Authorizations::help=AllSolAuthsHeader.html
solaris.account.:RO::Account Management::help=AccountHeader.html
<output omitted>
```

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

The slide shows an example of `auth_attr`, which is the configuration file that lists predefined authorization attributes. Each entry in the `auth_attr` database consists of one line of text containing six fields separated by colons (:), with the following format:

```
name:res1:res2:short_desc:long_desc:attr
```

Fields

- **name:** Unique string that is the name of the authorization
- **res1:** The characters `RO` indicate that this field is read-only and not modifiable by the tools that update this database.
- **res2:** Reserved for future use
- **short_description:** Short description or terse name for the authorization
- **long_description:** Reserved for future use
- **attr:** An optional list of semicolon-separated (;) key-value pairs that describe the attributes of the authorization. Zero or more keys can be specified. The keyword `help` identifies a `help` file in HTML.

Note: Authorizations can end with various suffixes:

- **.read** provides read access to user configuration files.
Example: `solaris.admin.usermgr.read`
- **.write** provides write access to user configuration files.
Example: `solaris.admin.usermgr.write`
- **.pswd** provides password access to user configuration files.
Example: `solaris.admin.usermgr.pswd`
- **.grant** permits a user to delegate to other users any assigned authorizations that begin with the same prefix.
Example: `solaris.admin.usermgr.grant`

Key RBAC Files: `exec_attr`

```
# getent exec_attr | grep 'Network Management'
Network
Management:solaris:cmd:RO::/usr/sbin/dladm:euid=dladm;egid=netadm;privs=
sys_dl_config,net_rawaccess,proc_audit
Network Management:solaris:cmd:RO::/usr/sbin/dlstat:euid=dladm;egid=sys
Network
Management:solaris:cmd:RO::/usr/sbin/flowadm:euid=dladm;egid=sys;privs=s
ys_dl_config,net_rawaccess,proc_audit
Network
Management:solaris:cmd:RO::/usr/sbin/flowstat:euid=dladm;egid=sys
Network
Management:solaris:cmd:RO::/usr/sbin/ipadm:euid=netadm;egid=netadm;privs
=sys_ip_config,net_rawaccess
Network Management:solaris:cmd:RO::/usr/bin/netstat:uid=0
Network Management:solaris:cmd:RO::/usr/bin/rup:euid=0
Network Management:solaris:cmd:RO::/usr/bin/ruptime:euid=0
Network Management:solaris:cmd:RO::/usr/bin/setuname:euid=0
Network Management:solaris:cmd:RO::/usr/sbin/asppp2pppd:euid=0
Network Management:solaris:cmd:RO::/usr/sbin/ifconfig:uid=0
...
<output truncated>
```

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

An execution attribute is associated with a rights profile name. An execution attribute can be a command with no options or a script that contains a command with options. Each entry in the `exec_attr` database consists of one line of text containing seven fields separated by colons (:), in the following basic format:

```
name:policy:type:res1:res2:id:attr
```

Fields

- **name:** Case-sensitive name of the profile
- **policy:** Security policy that is associated with the profile entry. The valid policies are `suser` (standard Oracle Solaris superuser) and `solaris`. The `solaris` policy recognizes privileges; the `suser` policy does not.
- **type:** Type of object defined in the profile. The `cmd` type specifies that the `ID` field is a command that is executed by a shell.

- **res1:** The characters `RO` in this field indicate that it is read-only and not modifiable by the tools that update this database.
- **res2:** Reserved for future use
- **id:** A string that uniquely identifies the object described by the profile. For a profile of type `cmd`, the ID is either the full path to the command or the asterisk (*) symbol, which is used to allow all commands. An asterisk that replaces the file name component in a path name indicates all files in a particular directory.
- **attr:** An optional list of semicolon-separated (;) key-value pairs that describe the security attributes to apply to the object upon execution. Zero or more keys can be specified. The list of valid keywords depends on the policy enforced. The following keywords are valid:
 - **euclid** and **uid:** Contain a single user name or a numeric user ID. Commands designated with `euclid` run with the effective UID indicated, which is similar to setting the `setuid` bit on an executable file. Commands designated with `uid` run with both real and effective UIDs.
 - **egid** and **gid:** Contain a single group name or a numeric group ID. Commands designated with `egid` run with the effective GID indicated, which is similar to setting the `setgid` bit on a file. Commands designated with `gid` run with both real and effective GIDs.
 - **privs:** Contains a privilege set that is added to the inheritable set before running the command
 - **Limitprivs:** Contains a privilege set that is assigned to the limit set before running the command

Note: `privs` and `limitprivs` are valid only for the `solaris` policy.

The example in the slide shows the commands and special security attributes for the `Printer Management` rights profile.

Key RBAC Files: prof_attr

```
# getent prof_attr | more
NTP Management:RO::Manage the NTP service:auths=solaris.smf.manage.ntp,solaris.s
mf.value.ntp
TPM Administration:RO::Administer Privileged TPM Operations:auths=solaris.smf.ma
nage.tcsd,solaris.smf.value.tcsd
D-BUS Management:RO::Manage D-BUS:auths=solaris.smf.manage.dbus;help=RtDBUSMngmn
t.html
DTrace Toolkit::::
Desktop Removable Media User:RO::Access removable media for desktop user:
Console User:RO::Manage System as the Console User:profiles=Desktop Removable Me
dia User,Suspend To RAM,Suspend To Disk,Brightness,CPU Power Management,Network
Autoconf User;auths=solaris.system.shutdown,solaris.device.cdrw,solaris.device.m
ount.removable,solaris.smf.manage.vbiosd,solaris.smf.value.vbiosd;help=RtConsUse
r.html
All:RO::Execute any command as the user or role:help=RtAll.html
Administrator Message Edit:RO::Update administrator message files:auths=solaris.
admin.edit/etc/issue,solaris.admin.edit/etc/motd;help=RtAdminMsg.html
Audit Configuration:RO::Configure Solaris Audit:auths=solaris.smf.value.audit;he
lp=RtAuditCfg.html
Audit Control:RO::Control Solaris Audit:auths=solaris.smf.manage.audit;help=RtAu
ditCtrl.html
Audit Review:RO::Review Solaris Auditing logs:help=RtAuditReview.html
Contract Observer:RO::Reliably observe any/all contract
events:help=RtContractObserver.html
<output omitted>
```

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

An execution profile is a mechanism that is used to bundle the commands and authorizations that are needed to perform a specific function. Each entry in the `prof_attr` database consists of one line of text containing five fields separated by colons (:), with the following format:

profname:res1:res2:desc:attr

Fields

- **name:** Case-sensitive name of the profile
- **res1:** The characters `RO` in this field indicate that it is read-only and not modifiable by the tools that update this database.
- **res2:** Reserved for future use
- **desc:** A long description that explains the purpose of the profile, including the type of user who would be interested in using it

- **attr:** An optional list of semicolon-separated (;) key-value pairs that describe the security attributes to apply to the object upon execution. There are four valid keys:
 - **help:** Is assigned the name of a file ending in .htm or .html
 - **auths:** Specifies a comma-separated list of authorization names chosen from the names defined in the `auth_attr` database. Authorization names can be specified by using the asterisk (*) character as a wildcard. For example, `solaris.printer.*` indicates all Oracle Solaris authorizations for printing.
 - **profiles:** Specifies a comma-separated list of profile names chosen from the names defined in the `prof_attr` database
 - **privs:** Specifies a comma-separated list of privilege names chosen from the names defined in the `priv_names` database

Relationships Among the RBAC Files

From the user_attr database:

```
sysadmin:::type=role;profiles=Device Management,File System Management,Printer
Management;roleauth=role

johndoe:::type=normal;auths=solaris.system.date;roles=sysadmin
```

From the prof_attr database:

```
Device Management:RO::Control Access to Removable
Media:auths=solaris.device.*;help=RtDeviceMngmnt.html
```

From the auth_attr database:

```
solaris.device.:RO::Device Allocation::help=DevAllocHeader.html
solaris.device.allocate:RO::Allocate Device::help=DevAllocate.html
solaris.device.config:RO::Configure Device Attributes::help=DevConfig.html
solaris.device.revoke:RO::Revoke or Reclaim Device::help=DevRevoke.html
solaris.device.cdrw:RO::CD-R/RW Recording Authorizations::help=DevCDRW.html
<output truncated>
```

From the exec_attr database:

```
Device Management:solaris:cmd:RO::/usr/sbin/allocate:uid=0
Device Management:solaris:cmd:RO::/usr/sbin/add_drv:uid=0
Device Management:solaris:cmd:RO::/usr/sbin/deallocate:uid=0
Device Management:solaris:cmd:RO::/usr/sbin/rem_drv:uid=0
Device Management:solaris:cmd:RO::/usr/sbin/update_drv:uid=0
```

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Now that you are familiar with the content of each of the four RBAC files, observe how the fields in the files are related.

The first section of the graphic shows a portion of a user_attr file. The user johndoe is a normal user account. The user is given the role of sysadmin, which is a role account. When assuming the sysadmin role, johndoe has access to specific rights profiles, which are defined as the Device Management, Filesystem Management, and Printer Management profiles.

From the sysadmin role entry in the first section to the next section (which is the prof_attr file), there is one relationship between the user_attr file and the prof_attr file. The Device Management rights profile, which is defined in the prof_attr file, is assigned to the sysadmin role in the user_attr file.

Observe the relationship between the `prof_attr` file and the `auth_attr` file, a portion of which is displayed in the third section of the graphic. The `Device Management` profile is defined in the `prof_attr` file as having all authorizations (beginning with the `solaris.device.*` string) assigned to it. These authorizations are defined in the `auth_attr` file.

Also observe the relationship between the `prof_attr` file and the `exec_attr` file, a portion of which is displayed in the fourth section. The `Device Management` profile is defined in the `prof_attr` file as having all authorizations (beginning with the `solaris.device.*` string) assigned to it.

Profile Shells

- Profile shells:
 - Enable access to the privileged rights that are assigned to the rights profile
 - Are assigned to a specific user as a login shell or through the `su` command to assume a role
- Users must be assigned one of the following profile shells:
 - `pfsh` for Bourne shell (`sh`)
 - `pfcsch` for C shell (`csh`)
 - `pfksh` for Korn shell (`ksh`)

The Oracle logo, consisting of the word "ORACLE" in a white, sans-serif font, centered within a red rectangular background.

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

A *profile shell* is a special type of shell that enables access to the privileged rights that are assigned to the rights profile. Standard UNIX shells cannot be used because they are not aware of the RBAC files and do not consult them.

Administrators can assign a profile shell to a specific user as a login shell, or the profile shell is started when that user runs the `su` command to assume a role.

Note: For descriptions of the profile shells, see the `pfexec(1)` man page.

When the user executes a command, the profile shell searches the role's rights profiles and associated rights. If the same command appears in more than one profile, the profile shell uses the first matching entry. The profile shell executes the command with the attributes that are specified in the RBAC configuration files.

Quiz

Which of the following files defines the rights profiles that are given to all new user accounts?

- a. /etc/shadow
- b. /etc/security/policy.conf
- c. usr/sbin/dladm/setuname

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Answer: b

Configuring RBAC

Configuring RBAC involves the following general actions:

- Creating a role
- Creating, cloning, or changing a rights profile
- Assigning a rights profile to a role
- Assigning a role to a user
- Assuming a role
- Restricting an administrator to explicitly assigned rights
- Assigning a rights profile to a user
- Delegating authorization to a user
- Assigning authorization to a role
- Modifying a system-wide RBAC policy

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Creating a Role

To create a role, use `roleadd -m -d dir rolename`.

```
# roleadd -u 3000 -g 10 -m -d /export/home/level1 -c "Level 1 Support" \
-P "Printer Management,Media Backup,Media Restore" level1
80 blocks
# passwd level1
New Password: <Type role password>
Re-enter new Password: <Type role password>
passwd: password successfully changed for level1
# getent passwd | grep level1
level1:x:3000:10:Level 1 Support:/export/home/level1:/bin/pfbash
# grep level1 /etc/shadow
level1:$5$j1ZvCrie$XJ6JRwBgUy0516/mNUQy596.h1R3f6g43gPDtd5v520:16027::::
::21808 # getent user_attr | grep level1
level1::::type=role;profiles=Printer Management,Media Backup,Media
Restore roleauth=role
```

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

To create a role, use the `roleadd` command combined with one or more options. The more common options include:

- **-u uid:** Specifies the user ID of the new role
- **-g gid:** Specifies an existing group's integer ID or character-string name
- **-m:** Creates the new role's home directory if it does not already exist
- **-d dir:** Specifies the home directory of the new role
- **-c comment:** Text string that provides a short description of the role
- **-P profile:** Assigns rights profiles to the role. Use commas (,) to separate multiple rights profiles.
- **rolename:** Name of the new role. For restrictions on acceptable strings, see the `roleadd (1M)` man page.

Note: To create a role, you must be an administrator with the `User Management` rights profile. To assign a password to the role, you must be assigned the `User Security` rights profile.

The `roleadd` command creates a role entry in the `/etc/passwd`, `/etc/shadow`, and `user_attr` files. In this example, the `roleadd` command creates a new role called `level1`, builds the home directory, and assigns the role with rights profiles of `Printer Management`, `Media Backup`, and `Media Restore` to the user ID 3000 and group ID 10. The role cannot be used until a password is applied to it.

Note: The installation of the Oracle Solaris 11 OS has the `Printer Management`, `Media Backup`, and `Media Restore` rights profiles already defined in the `exec_attr` and `prof_attr` files, so there is no need to add an entry for these profiles in these two files. However, if you want to create a new rights profile, you must make a new entry in the `prof_attr` file. You learn how to do that next.

The changes to the `/etc/passwd`, `/etc/shadow`, and `user_attr` files are shown in the example. The type of this account is `role (type=role)` and includes the `Printer Management`, `Media Backup`, and `Media Restore` rights profiles.

Creating a Rights Profile

To create a profile, use `profiles -p -S "Profile"`.

```
# profiles -p -S LDAP "Test Users"
profiles:Sun Ray Users> set desc="For all users"
profiles:Sun Ray Users> add profiles="Test Basic User"
profiles:Sun Ray Users> set defaultpriv="basic,!proc_info"
profiles:Sun Ray Users> set limitpriv="basic,!proc_info"
profiles:Sun Ray Users> end ... Ray Users> exit
#
# profiles -p "Test Users"
Found profile in LDAP repository.
profiles:Sun Ray Users> info
    name=Test Users
    desc=For all users
    defaultpriv=basic,!proc_info,
    limitpriv=basic,!proc_info,
    profiles=Test Basic User
```



Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Use the `set` subcommand for profile properties that have a single value (such as `set desc`) and the `add` subcommand for properties that have more than one value (such as `add profiles`).

In this example, the administrator creates a rights profile for Test Users in the LDAP repository. The administrator has already created a test version of the Basic Solaris User rights profile and has removed all rights profiles from the `policy.conf` file on the Test server. The administrator verifies the content with the `info` command.

Cloning a Rights Profile

- The rights profiles that Oracle Solaris provides are read-only.
- You can clone a provided rights profile for modification if its collection of security attributes is insufficient.
- You can then continue to modify the new rights profile by adding or removing supplementary rights profiles, authorizations, and other security attributes.

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

For example, you might want to add the `solaris.admin.edit/path-to-system-file` authorization to a provided rights profile.

Modifying a Rights Profile

You can modify a rights profile to do the following:

- Enhance an existing rights profile
 1. Create a new profile.
 2. Add the existing rights profile as a supplementary rights profile.
 3. Add the enhancements.
- Remove content from an existing rights profile
 1. Clone the profile.
 2. Rename the profile.
 3. Modify the profile.

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Modifying a Rights Profile

```
# profiles -p "Network IPsec Management"
profiles:Network IPsec Management> add auths="solaris.admin.edit/etc/hosts"
Cannot add. Profile cannot be modified
#
# profiles -p "Total IPsec Mgt"
Total IPsec Mgt> set desc="Network IPsec Mgt plus edit authorization"
Total IPsec Mgt> add profiles="Network IPsec Management"
Total IPsec Mgt> add auths="solaris.admin.edit/etc/hosts"
Total IPsec Mgt> add auths="solaris.admin.edit/etc/inet/ipsecinit.conf"
Total IPsec Mgt> add auths="solaris.admin.edit/etc/inet/ike/config"
Total IPsec Mgt> add auths="solaris.admin.edit/etc/inet/secret/ipseckeys"
Total IPsec Mgt> end
Total IPsec Mgt> exit
#
# profiles -p "Total IPsec Mgt" info
      name=Total IPsec Mgt
      desc=Network IPsec Mgt plus edit authorization
      auths=solaris.admin.edit/etc/hosts,
            solaris.admin.edit/etc/inet/ipsecinit.conf,
            solaris.admin.edit/etc/inet/ike/config,
            solaris.admin.edit/etc/inet/secret/ipseckeys
      profiles=Network IPsec Management
```



Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

In this example, the administrator adds several `solaris.admin.edit` authorizations to a site IPsec Management rights profile. The administrator verifies that the Network IPsec Management rights profile cannot be modified. The administrator then creates a rights profile that includes the Network IPsec Management profile. The administrator verifies the content.

Assigning a Rights Profile to a Role

To assign a rights profile to a role, use `rolemod [-P profile] [-s shell] rolename`.

```
# rolemod -P profile1,profile2 -s /usr/bin/pfksh level1
```

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

The `rolemod` command changes the definition of the specified role and makes the appropriate login-related changes to the system file and file system. Note that the `rolemod` command modifies the entry for the specified role in the `/etc/passwd`, `/etc/shadow`, and `user_attr` files.

You can use the following options with the `rolemod` command:

- **-e expire:** The date on which a role expires. Use this option to create temporary roles.
- **-l new_logname:** Specifies the new login name for the role
- **-P profile:** Specifies one or more comma-separated rights profiles, as defined in the `prof_attr` file
- **-s shell:** Login shell for `rolename`. This shell must be a profile shell.
- **rolename:** Name of the role you are modifying

In the example in the slide, the `profile1` and `profile2` profiles and the `/usr/bin/pfksh` profile shell are assigned to the role named `level1`.

Assigning a Role to a User

- A user can have access to many roles.
- To add a role to a new user:
 1. Use the `useradd` command.
 2. Assign a password to the role by using `passwd rolename`.
 3. Verify that an entry has been made in the `user_attr` file.
- To add roles to an existing user account, use `usermod`.

The Oracle logo, consisting of the word "ORACLE" in a white, sans-serif font, centered within a solid red rectangular background.

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Notes for step 2: If you are assigned the *User Security* rights profile, you can create the password. Otherwise, a user who is assigned the role must create the password by using the `su - rolename` command. Because a role account is assigned to more than one user, the superuser usually creates a role password and provides users with that role.

Note: To remove all role access from a user account, use the `usermod` command with the `-R ""` option followed by the user login name.

Assigning a Role to a User

```
# useradd -u 4009 -g 10 -m -d /export/home/paul \
-R level1 -c "Paul" paul
80 blocks
# passwd paul
New Password: <Type rolename password>
Re-enter new Password: <Type rolename password>
passwd: password successfully changed for paul
# getent user_attr | grep paul
paul:::roles=level1
# roles paul
level1
# usermod -R level1 chris
# passwd -r repository level1
Password: <Type rolename password>
Confirm Password: <Retype rolename password>
# usermod -R "" chris
```

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

The example in the slide shows the `useradd` command used with the `-R` option to define the `level1` role for the user `paul`. To verify that the `level1` role is assigned to `paul`, first view the `user_attr` file for the user `paul`. Observe that the entry for `paul` has the `level1` role. You can also use the `roles` command to see the roles that are assigned to `paul`.

Note: The association between the `paul` user account and the `level1` role is defined automatically in the `user_attr` file.

The `level1` role is then assigned to the existing user account `chris` by using the `usermod -R` command. In the last line, all role access is removed from the `chris` account by using the `usermod -R ""` command.

Assuming a Role

To assume a role:

1. Use `roles` to determine the roles that you can assume.
2. Use `su - rolename` to assume a role.
3. Verify that you are now in a role by using `/usr/ucb/whoami`.
4. View the capabilities of your role by using `ppriv $$`.

```
# roles
sysadmin,oper,primaryadm
# su - sysadmin
Password: <Type sysadmin password>
$/usr/bin/whoami
sysadmin
$ ppriv $$
950:    bash
flags = <none>
      E: basic
      I: basic
      P: basic
      L: all
```

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Notes for step 4: In contrast to the `root` role, the `System Administrator` role has the basic set of privileges in its effective (E) set.

In the example in the slide, the capabilities are all `basic` except for the limit (L) privilege set, which by default has the capability `all`.

Restricting an Administrator to Explicitly Assigned Rights

You can restrict a role or a user to a limited number of administrative actions in two ways:

- Use the `Stop` rights profile.
- Modify the `policy.conf` file on a system, and require the role or user to use that system for administrative tasks.

```
# rolemod -P "Profile_Name,All,Stop" rolename
```

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

The `Stop` rights profile is the simplest way to create a restricted shell. The authorizations and rights profiles that are assigned in the `policy.conf` file are not consulted. In the default configuration, the role or user is not assigned the `Basic Solaris User` rights profile, the `Console User` rights profile, or the `solaris.device.cdrw` authorization.

The `rolemod -P` command is used with the `Stop` rights profile, as shown in the code example in the slide. This command is especially useful if you have many profiles assigned to a role and you want to limit the role to only a few profiles.

Assigning the Rights Profile to a User

- To assign a rights profile to a user, use `usermod`.
- The `usermod` command automatically updates the `user_attr` file for the specified user.

```
# profiles chris
Basic Solaris User
All
# usermod -P "Printer Management" chris
# profiles chris
Printer Management
Basic Solaris User
All
# getent user_attr | grep chris
chris:::profiles=Printer Management;roles=
# profiles -l chris
Printer Management:
/etc/init.d/lp euid=0, uid=0
/usr/bin/cancel euid=lp, uid=lp
<output omitted>
All:
*
```

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

The rights profiles that are assigned to a user can be listed with the `profiles` command. Every account has the `All` rights profile, which allows any command to be executed (but with special security attributes).

Note: Other rights profiles that are assigned to all new user accounts are defined in the `/etc/security/policy.conf` file.

The code example in the slide shows the `Printer Management` rights profile assigned to the `chris` user account. If you run the `profiles` command again for the user, you can see that the `Printer Management` rights profile has been added.

The `usermod` command automatically updates the `user_attr` file for the specified user, as shown in the code example. The new line for the user `chris` shows the new profile assignment.

You can examine the content of a rights profile with the `-l` option of the `profiles` command. The individual commands in the rights profile can be seen, along with the special security attributes with which they are executed.

Delegating an Authorization to a User

- Authorizations can be assigned to user accounts.
- Authorizations can also be assigned to roles or embedded in a rights profile that can be assigned to a user or role.

To delegate an authorization to a user:

1. Use `usermod -A authorization loginname`.
2. Verify that an entry has been made in the `user_attr` file for the user.
3. View the authorizations for the user by using `auths`.

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

To delegate an authorization to a user, use the `usermod` command with the `-A` option, the authorization, and the user login name.

Note: Only a user or role that has grant rights to an authorization can assign the authorization to an account. The `roleadd` command automatically updates the `user_attr` file.

To verify that the authorization has been assigned to the user, check the `user_attr` file. You can also use the `auths` command for the user to see if the authorization is listed in the entry.

Delegating an Authorization to a User

```
# su - chris
Oracle Corporation      SunOS 5.11  11.1      September 2012
chris:~$ crontab -l root
crontab: you must be super-user to access another user's crontab file
chris:~$ exit
# usermod -A solaris.jobs.admin chris
# getent user_attr | grep chris
chris:::auths=solaris.jobs.admin;profiles=Printer Management,roles=
# auths chris
solaris.admin.printer.read,solaris.admin.printer.modify,solaris.admin.pri
nter.delete,solaris.device.cdrw,solaris.profmgr.read,solaris.jobs.users,s
olaris.mail.mailq,solaris.admin.usermgr.read,solaris.admin.logsvc.read,so
laris.admin.fsmgr.read,solaris.admin.serialmgr.read,solaris.admin.diskmgr
.read,solaris.admin.procmgr.user,solaris.compsys.read,solaris.admin.prodr
eg.read,solaris.admin.dcmgr.read,solaris.snmp.read,solaris.project.read,s
olaris.admin.patchmgr.read,solaris.network.hosts.read,solaris.admin.volmg
r.read
# su - chris
Oracle Corporation      SunOS 5.11  11.1      September 2011
chris:~$ crontab -l root
#ident "%Z%M%  %I% %E% SMI"
#
# The root crontab should be used to perform accounting data collection.
(output omitted)
chris:~$ exit
```



Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

In this example, a regular user is not permitted to look at another user's crontab file. To grant the user authorization to manage the other user's crontab file, use the `usermod` command with the `-A` option to add an authorization. The `user_attr` file is automatically modified with this new information.

The `chris` account is a normal user account (`type=normal`). You can see in the `user_attr` file that the `solaris.jobs.admin` authorization and the Printer Management rights profile were added previously for `chris`. You then use the `auths` command to see the authorizations assigned to `chris`. With this authorization, `chris` can now view or modify other users' crontab files.

Assigning Authorization to a Role

You can assign authorizations to a role and then give users access to that role.

To assign authorization to a role:

1. Use `rolemod -A "authorization" <rolename>`.
2. Verify that an entry has been made in the `user_attr` file for the role.
3. View the authorizations for the role by using `auths`.

```
# rolemod -A "solaris.admin.usermgr.*" level2
# auths level2
solaris.admin.usermgr.*
```

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

If a large number of user accounts require the same configuration and management of authorizations, it is easier to assign the authorizations to a role and give users access to that role. You can assign the authorization to a role by using the `rolemod -A` command. Note that the `rolemod` command automatically updates the `user_attr` file.

In the example in the slide, the `solaris.admin.usermgr.*` authorization is assigned to the `level2` role. The `auths` command is used to verify that the authorization has been assigned to the role.

Rights Profiles

Rights profiles that are given to all new user accounts are defined in the `/etc/security/policy.conf` file.

```
# cat /etc/security/policy.conf
<header and copyright output omitted>
#
AUTHS_GRANTED=solaris.device.cdrw
PROFS_GRANTED=Basic Solaris User
CONSOLE_USER=Console User
# crypt(3c) Algorithms Configuration
#
# CRYPT_ALGORITHMS_ALLOW specifies the algorithms that are allowed to
# be used for new passwords. This is enforced only in crypt_gensalt(3c).
#
CRYPT_ALGORITHMS_ALLOW=1,2a,md5,5,6
<output omitted>
#PRIV_DEFAULT=basic
#PRIV_LIMIT=all
<output truncated>...
```



Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

The settings in the `/etc/security/policy.conf` file determine the default privileges that users have. If these settings are not set, the default privileges are taken from the inherited set.

Settings

- **PRIV_DEFAULT:** Determines the default set on login
- **PRIV_LIMIT:** Defines the Limit set on login. Individual users can have privileges assigned or taken away by the use of `user_attr`.

Modifying a System-Wide RBAC Policy

- The `/etc/security/policy.conf` file establishes a system-wide RBAC policy.
- There are two different settings for the system-wide policy:
 - `PRIV_DEFAULT`: Determines the default
 - `PRIV_DEFAULT=basic,!file_link_any`: Can be used to modify the default
- The default is set to `PRIV_DEFAULT=basic`.
- You can modify this file to deny specific privileges to non-administrative users.

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Modifying a System-Wide RBAC Policy

To modify the system-wide policy:

1. Change the `PRIV_DEFAULT=basic` default entry by using a text editor.
2. Reboot the system so that the changes take effect.
3. Test the modification as a user.

```
# vi /etc/security/policy.conf
# grep PRIV_DEFAULT /etc/security/policy.conf
# There are two different settings; PRIV_DEFAULT determines the default
# Similarly, PRIV_DEFAULT=basic,!file_link_any takes away only the
PRIV_DEFAULT=basic,!proc_info,!proc_session
# init 6
<log in to the system>
# su - jjones
Oracle Corporation SunOS 5.11      11.1      September 2012
jjones:~$ ps -A -o user -o pid -o comm | more
      USER  PID COMMAND
jjones 1941  ps
jjones 1935 -bash
```

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

The example in the slide shows how to deny a non-administrative user the privilege to look at the processes of other users. Edit the `PRIV_DEFAULT=basic` entry as follows:

```
PRIV_DEFAULT=basic, !proc_info, !proc_session
```

For the changes to take place, reboot the system. After logging back in to the system, `su` to the `jjones` user account and issue the command to access the processes. The only processes that the user can display are the user's own processes.

Note: The `-A` and `-o` options that are used in the `ps` command tell the system to write information for all processes in the specified format. In the example, this format is by `user`, `pid`, and `command`.

Quiz

Which of the following rights profiles must an administrator have to create a role?

- a. User Security
- b. Basic Solaris User
- c. User Management
- d. IPsec Management

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Answer: c

Summary

In this lesson, you should have learned how to:

- Administer process rights management
- Configure RBAC


ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Installing the Oracle Solaris 11 Operating System

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.



System Administration for Experienced UNIX/Linux Administrators



Administering System
Software by Using IPS



Administering Services
by Using SMF



Administering ZFS



Configuring the Network



Administering Oracle Solaris
Zones



Administering Privileges
and RBAC



Installing the Oracle Solaris 11
Operating System



Monitoring System Resources

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Objectives

After completing this lesson, you should be able to:

- Perform pre-installation tasks
- Install Oracle Solaris 11 on a single host by using the Live Media Installer
- Install Oracle Solaris 11 on a single host by using the Text Installer
- Install Oracle Solaris 11 on multiple hosts by using the Automated Installer

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Lesson Agenda

- Performing pre-installation tasks
- Installing Oracle Solaris 11 on a single host by using the Live Media Installer
- Installing Oracle Solaris 11 on a single host by using the Text Installer
- Installing Oracle Solaris 11 on multiple hosts by using the Automated Installer

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Oracle Solaris Installation





- Oracle Solaris is an enterprise-class operating system (OS) that enables customers to build datacenter-grade infrastructures on a wide range of SPARC and x86 servers and on Oracle engineered systems.
- Installing the Oracle Solaris 11 OS consists of the following phases:
 1. Preparing for the installation
 2. Performing the installation
 3. Verifying the installation
 4. Rebooting the system

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Preparing for the Installation

Preparing for an installation includes addressing the following pre-installation tasks:

	Reviewing the Oracle Solaris 11.1 Release Notes
	Selecting the installation option
	Identifying system requirements
	Downloading the ISO image

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Reviewing the Release Notes

- An important exercise while preparing for an installation is to review the Oracle Solaris 11.1 [Release Notes](#).
- The Release Notes describe important installation, update, and runtime issues that you might need to consider before installing or running Oracle Solaris 11.1.

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Selecting the Installation Option

iso Images	x86	SPARC	Single Host	Multiple Hosts
Live Media	✓			
Text Installer	✓	✓	✓	
Automated Installer (AI)	✓	✓	✓	✓

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

With Oracle Solaris 11, the installer on the Live Media ISO image is for 64-bit x86 platforms *only*. If you must install the operating system on a SPARC-based system (M-Series and T-Series servers only), you must use the Text Installer or the Automated Installer.

Interactive installers can perform an initial installation on:

- Whole disk
- Oracle Solaris x86 partition
- SPARC slice (Text Installer)

Caution: The installation process overwrites any data that exists on the disk that you have identified or targeted for the installation.

Identifying System Requirements

System Component	Requirement
Platform	<ul style="list-style-type: none"> • M-Series (+ Dynamic Domains) • T-Series (+ Oracle VM for SPARC, formerly known as “LDoms”) • x86 (64-bit processor) (+ Oracle VM for x86)
Disk space	Recommended minimum: 13 GB
Memory	Recommended minimum: <ul style="list-style-type: none"> • Text Installer: 1 GB • Live Media: 1.5 GB

Note: The recommended minimum values are subject to change with the final release of the software. See the release notes for final disk space and memory recommendations.

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

If the GUI installer on the Live Media ISO image does not work on your system, use the Text Installer.

Note: SPARC support is available on M-Series and T-Series systems only. OpenBoot PROM (OBP) is required to be at 4.17 or higher. Using the latest firmware is recommended.

Downloading Images

- To install the OS on an x86 system by using the Live Media, Text Installer, or AI options, download the following ISO images:
 - [Download for x86 - Live Media](#)
 - [Download for x86 - Text Install](#)
 - [Download for x86 - AI](#)
- To install the OS on a SPARC system by using the Text Installer or AI options, download the following ISO images:
 - [Download for SPARC - Text Install](#)
 - [Download for SPARC - AI](#)

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Agenda

- Performing pre-installation tasks prior to installing Oracle Solaris 11
- Installing Oracle Solaris 11 on a single host by using the Live Media Installer
- Installing Oracle Solaris 11 on a single host by using the Text Installer
- Installing Oracle Solaris 11 on multiple hosts by using the Automated Installer

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Installing Oracle Solaris 11 by Using the Live Media Installer

```
USB keyboard
1. Arabic
2. Belgian
3. Brazilian
4. Canadian-Bilingual
5. Canadian-French
6. Danish
7. Dutch
8. Dvorak
9. Finnish
10. French
11. German
12. Italian
13. Japanese-type6
14. Japanese
15. Korean
16. Latin-American
17. Norwegian
18. Portuguese
19. Russian
20. Spanish
21. Swedish
22. Swiss-French
23. Swiss-German
24. Traditional-Chinese
25. TurkishQ
26. UK-English
27. US-English
To select the keyboard layout, enter a number [default 27]:
```



Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

The Oracle Solaris 11 Live Media Installer provides a GUI (graphical user interface) that walks you through the process of configuring the system for the initial OS installation. The ISO image boots to a full OS with a functional desktop. After you boot the Live Media with the GUI installer, the first thing you are asked to do is identify the keyboard layout. The default is US-English [27].

Installing Oracle Solaris 11 by Using the Live Media Installer

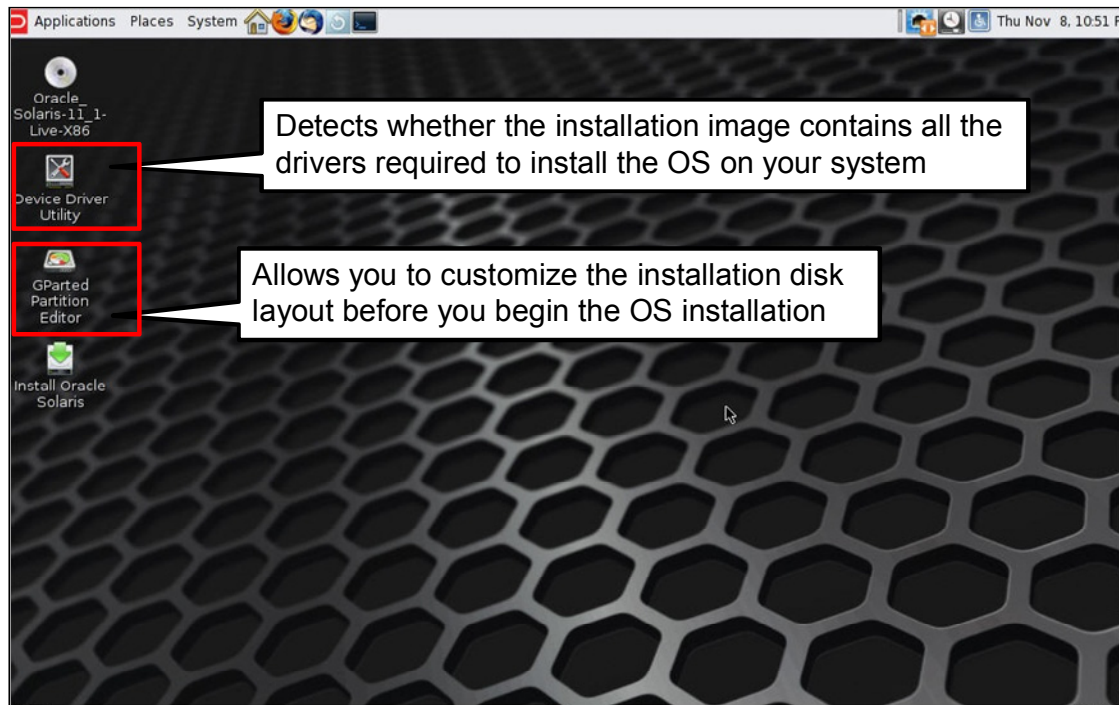
```
1. Chinese - Simplified
2. Chinese - Traditional
3. English
4. French
5. German
6. Italian
7. Japanese
8. Korean
9. Portuguese - Brazil
10. Spanish
To select the language you wish to use, enter a number [default is 3]: █
```

The Oracle logo, consisting of the word "ORACLE" in a white, sans-serif font, is centered on a solid red horizontal bar.

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

After selecting the keyboard layout, you are prompted to select the language you want to use. Again, English is the default [3]. From this point, the installer configures the system devices and then launches the GUI.

Introducing the Live Media Desktop



ORACLE

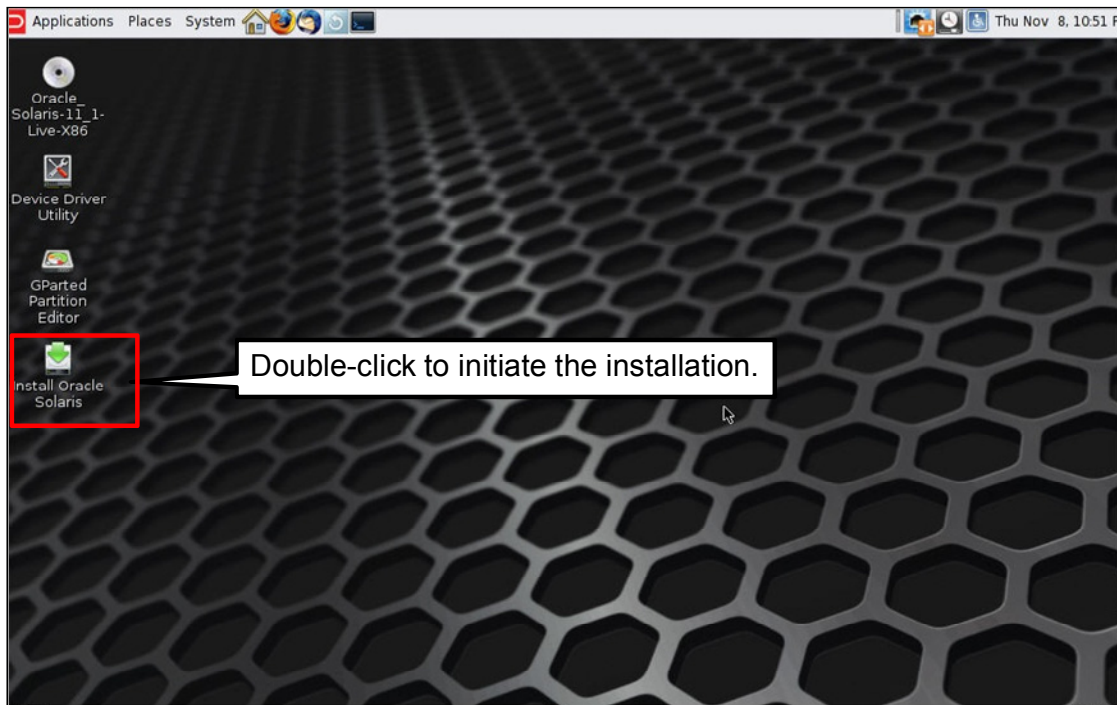
Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

The GUI provides a full desktop operating environment (as shown in this slide). The Live Media Installer provides additional tools to assist you in your installation, such as the Device Driver Utility and a partition editor.

When you boot your system from the Live Media ISO image and see the desktop, the Device Driver Utility automatically launches and begins searching for missing device drivers. If the utility locates any such devices, a notification is displayed on the desktop.

You also have the option of manually launching the Device Driver Utility by double-clicking the icon or selecting Applications > System > Device Driver Utility.

Initiating the Installation with Live Media

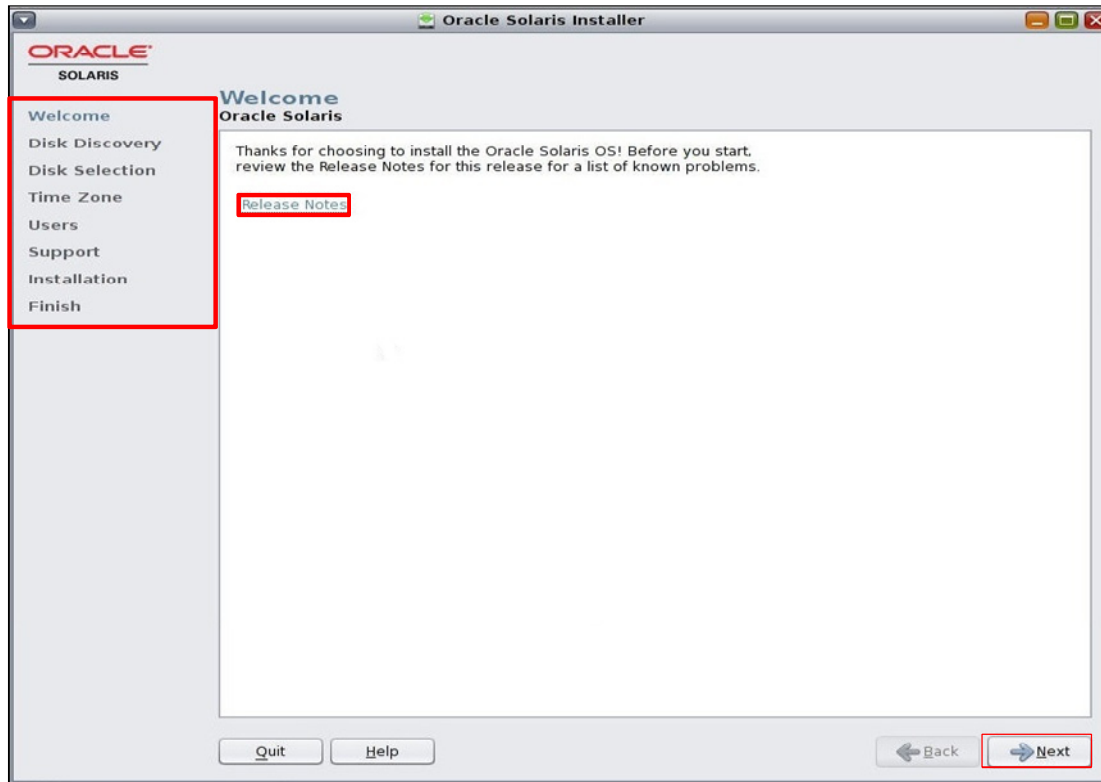


ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

To initiate the installation, double-click the **Install Oracle Solaris** icon.

Welcome Screen




Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

The first screen you see is the Welcome screen. From this screen, you can review the Release Notes (if you have not already done so).

The list of items on the left of the Welcome screen highlights the steps to take to complete the installation. You begin the installation by providing configuration data for the disk, time zone, and users.

After you provide the required information in each of the configuration data screens, the actual installation begins.

Oracle Solaris 11 Live Media: Disk Discovery



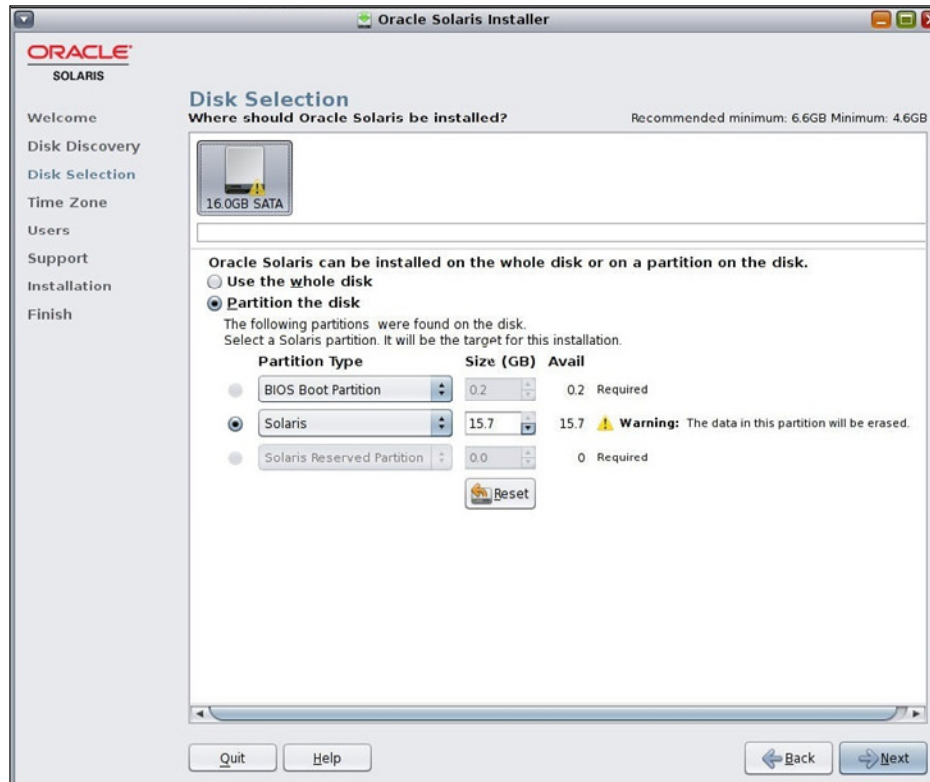
ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Starting with the Oracle Solaris 11.1 release, the ability to install to iSCSI target Logical Unit Numbers (LUNs) has been included in the Live Media Installer. Administrators can choose between installing on local disks or iSCSI disks.

- **Local Disks:** This is the default option for disks that are attached to the computer, including internal and external hard disks.
- **iSCSI:** Select this option if you want the installer to search for remote disks that are accessible over a network using the iSCSI standard. You can connect to a remote iSCSI disk by using DHCP auto-discovery or by manually specifying a target IP address, iSCSI target name and LUN, and initiator name.

Selecting a Disk



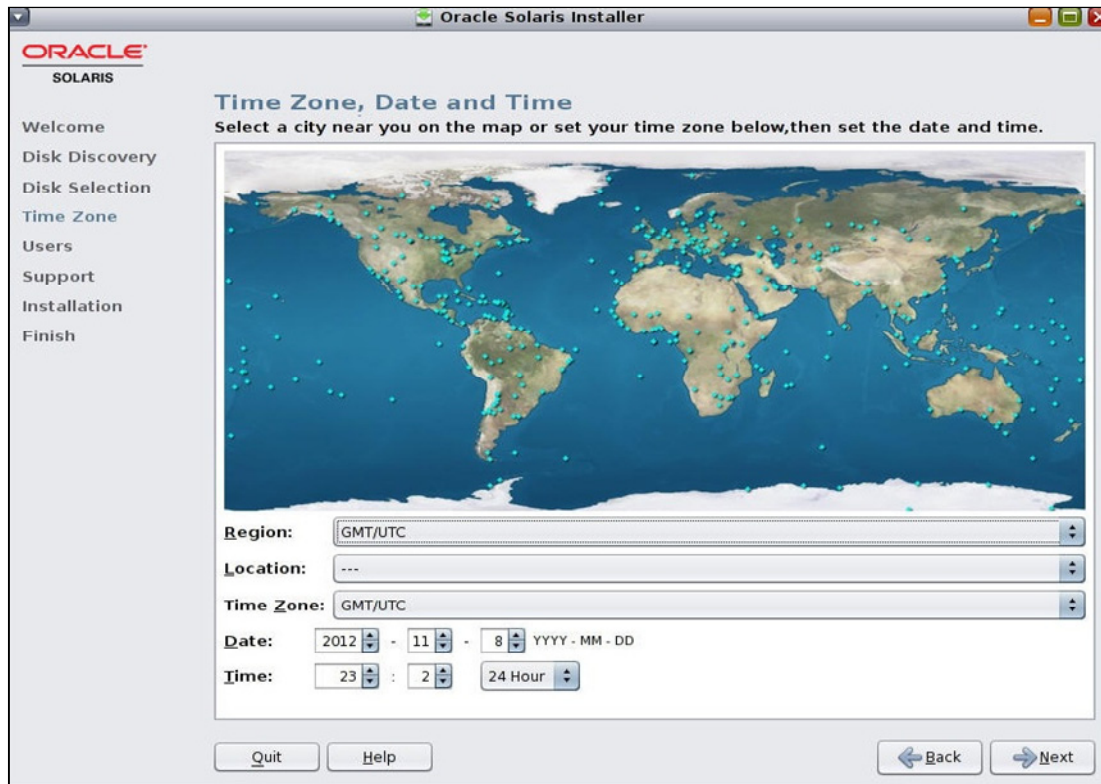
ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

On the Disk screen, you are prompted to select where you want Oracle Solaris to be installed:

- **Use the whole disk**
- **Partition the disk:** If you select a disk partition, you must also select the partition type and size.

Setting the Time Zone, Date, and Time



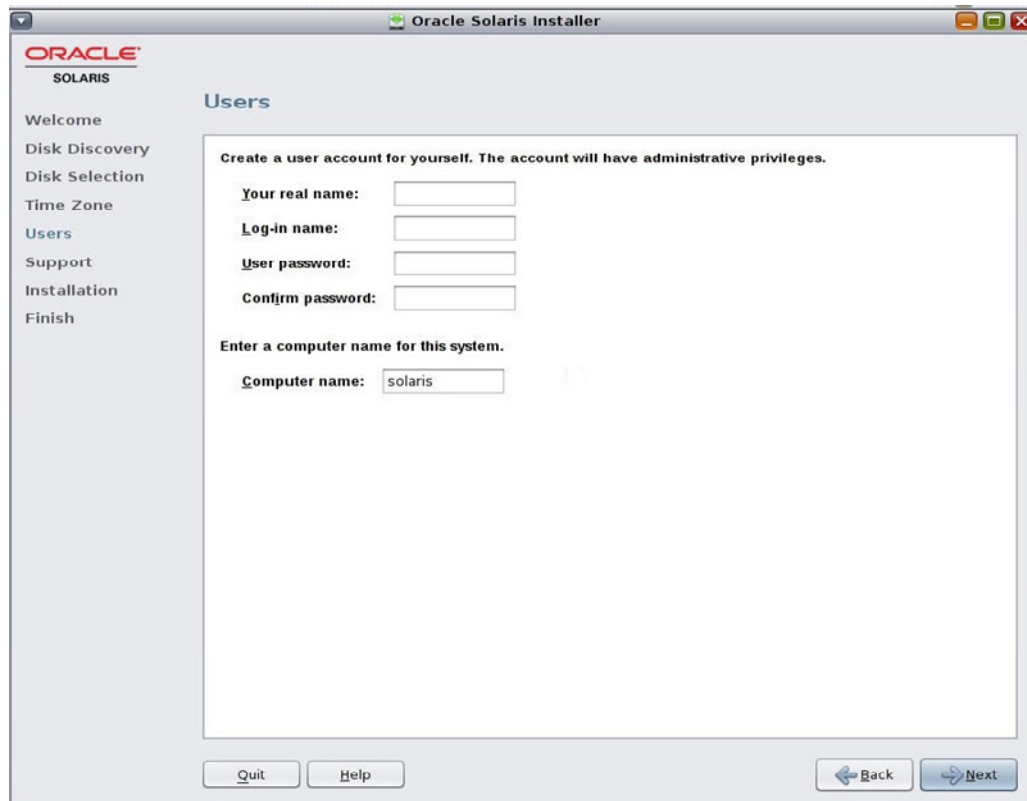
ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

The next screen that appears is “Time Zone, Date and Time,” where you can select the region, location, and time zone that are appropriate to your installation. You can also set the date and time.

To continue, click **Next**.

Providing User Information

**ORACLE**

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

On the **Users** screen, you enter the user information, including real name, login name, and user password. You are also asked to provide a computer name, which is also referred to as the *host name* or *node name*.

To continue, click **Next**.

Note: The first configured user is given the `root` role.

Support Registration

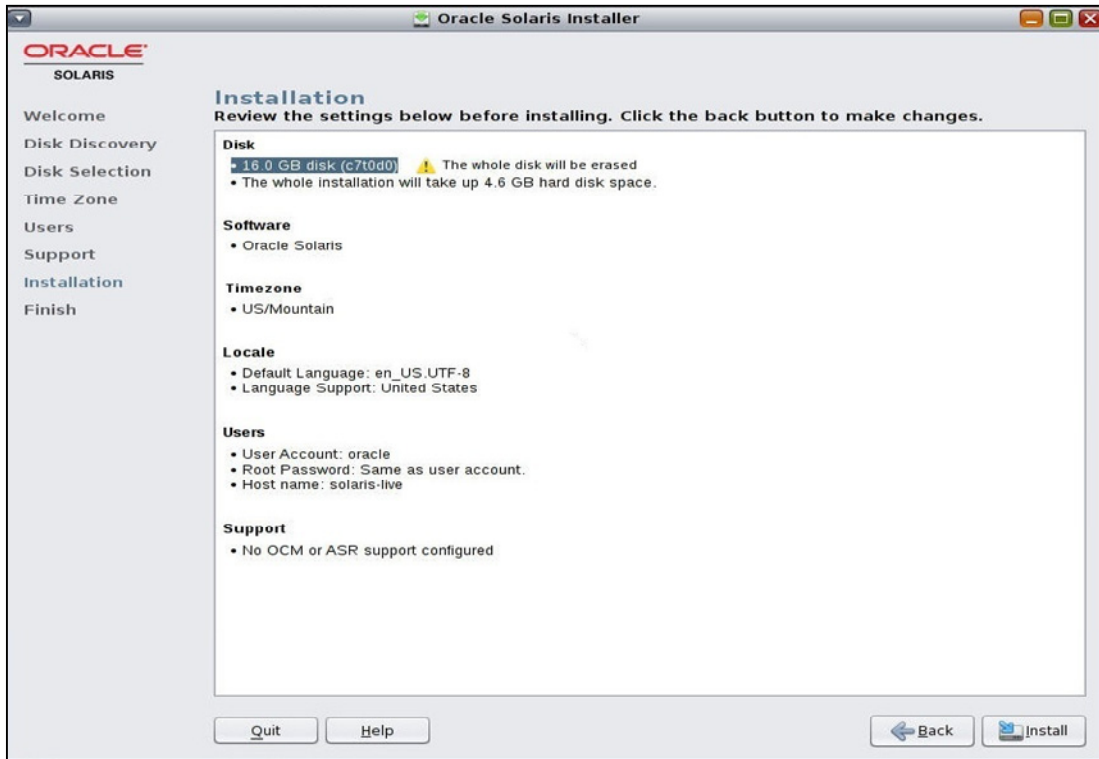
The screenshot shows the 'Oracle Solaris Installer' window with the 'Support Registration' tab selected. On the left is a navigation pane with options: Welcome, Disk Discovery, Disk Selection, Time Zone, Users, Support (highlighted), Installation, and Finish. The main content area is titled 'Support Registration' and contains the following text: 'Provide your email address to be informed of security issues, install the product and initiate configuration manager. Please see <http://www.oracle.com/goto/solarisautoreg> for details.' Below this is an 'Email:' field with 'anonymous@oracle.com' entered. A note says 'Easier for you if you use your My Oracle Support email address/username.' There is a checked checkbox for 'I wish to receive security updates via My Oracle Support.' and a 'My Oracle Support Password:' field. At the bottom, under 'Network for support registration:', there are three radio button options: 'No proxy - direct connect to the internet' (selected), 'Proxy - configure proxy information', and 'Aggregation Hubs - specify hubs used for centralized support data'. At the very bottom of the window are 'Quit', 'Help', 'Back', and 'Next' buttons.

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

On the Support Registration screen, you provide your email address and password. You can also select the option that enables you to receive security updates via My Oracle Support. The “No proxy” option is selected by default.

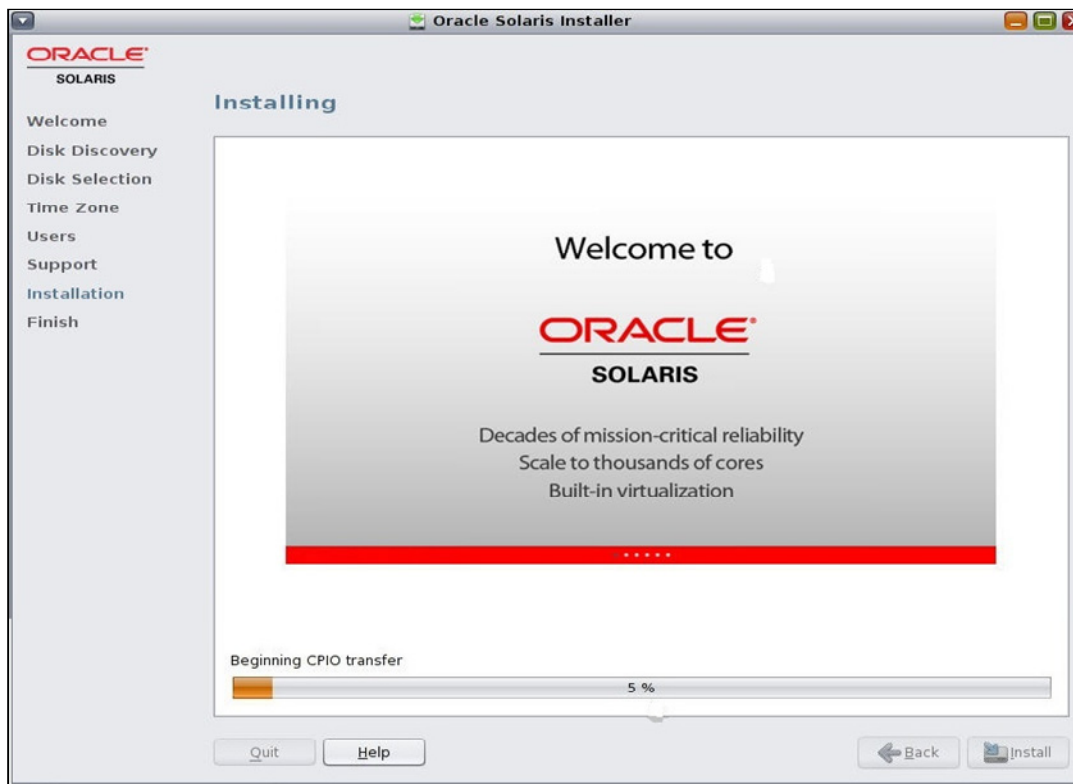
Reviewing Installation Specifications

**ORACLE**

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

After completing the configuration data, you see the Installation screen. Review the information carefully to ensure that it is accurate before you start the installation. You can go back and make changes later if necessary.

Monitoring the Installation



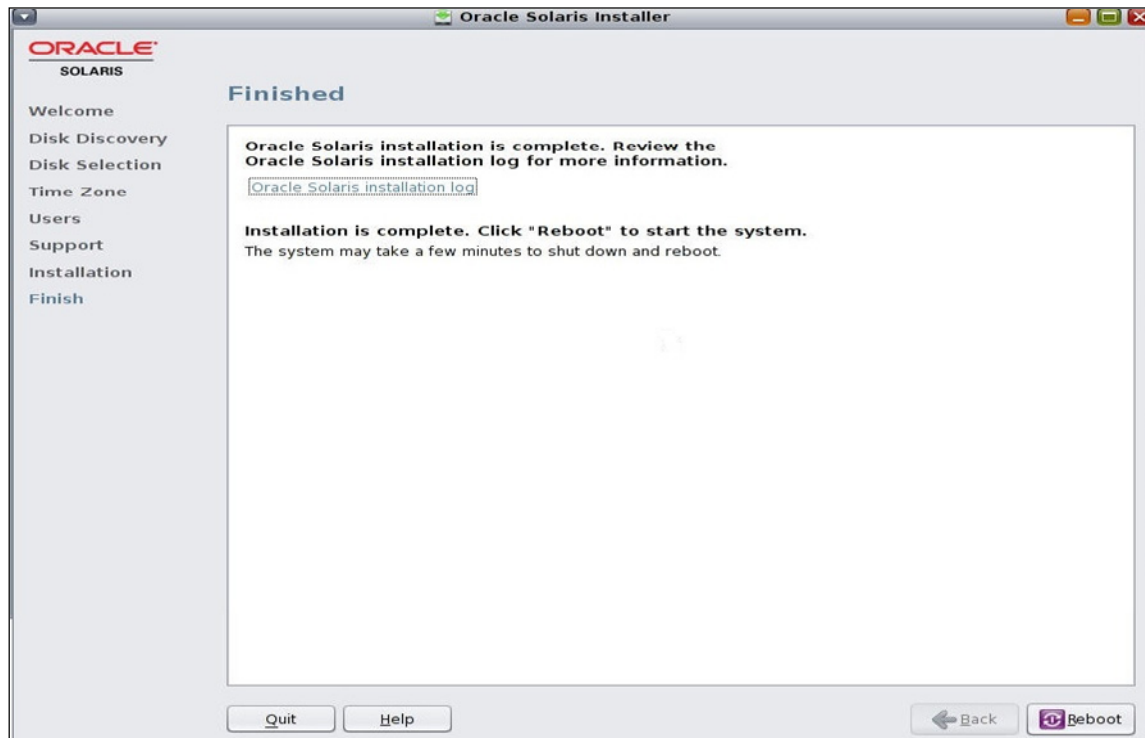
ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

The Installation screen enables you to monitor the progress of the installation. The installation takes about 15 to 20 minutes to complete.

Caution: After the installation starts, do not interrupt it. Interrupting an installation can leave a disk in an indeterminate state.

Verifying the Installation



ORACLE

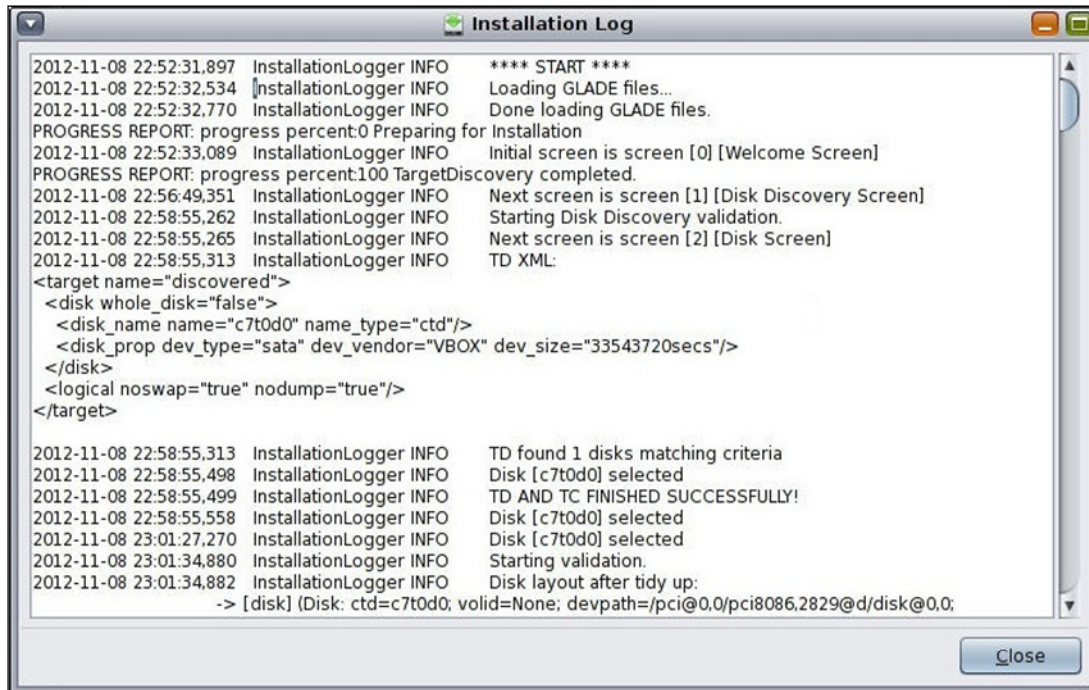
Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

When the installation concludes, you see the Finished screen, which provides you with access to the installation log and an opportunity to verify that:

- No errors occurred during installation
- Major facilities were successfully installed

To access the log, click the **Oracle Solaris installation log** link. A separate dialog box then appears with the log content.

Reviewing the Installation Log



```

2012-11-08 22:52:31.897 InstallationLogger INFO **** START ****
2012-11-08 22:52:32.534 InstallationLogger INFO Loading GLADE files...
2012-11-08 22:52:32.770 InstallationLogger INFO Done loading GLADE files.
PROGRESS REPORT: progress percent:0 Preparing for Installation
2012-11-08 22:52:33.089 InstallationLogger INFO Initial screen is screen [0] [Welcome Screen]
PROGRESS REPORT: progress percent:100 TargetDiscovery completed.
2012-11-08 22:56:49.351 InstallationLogger INFO Next screen is screen [1] [Disk Discovery Screen]
2012-11-08 22:58:55.262 InstallationLogger INFO Starting Disk Discovery validation.
2012-11-08 22:58:55.265 InstallationLogger INFO Next screen is screen [2] [Disk Screen]
2012-11-08 22:58:55.313 InstallationLogger INFO TD XML:
<target name="discovered">
  <disk whole_disk="false">
    <disk_name name="c7t0d0" name_type="ctd"/>
    <disk_prop dev_type="sata" dev_vendor="VBOX" dev_size="33543720secs"/>
  </disk>
  <logical noswap="true" nodump="true"/>
</target>

2012-11-08 22:58:55.313 InstallationLogger INFO TD found 1 disks matching criteria
2012-11-08 22:58:55.498 InstallationLogger INFO Disk [c7t0d0] selected
2012-11-08 22:58:55.499 InstallationLogger INFO TD AND TC FINISHED SUCCESSFULLY!
2012-11-08 22:58:55.558 InstallationLogger INFO Disk [c7t0d0] selected
2012-11-08 23:01:27.270 InstallationLogger INFO Disk [c7t0d0] selected
2012-11-08 23:01:34.880 InstallationLogger INFO Starting validation.
2012-11-08 23:01:34.882 InstallationLogger INFO Disk layout after tidy up:
-> [disk] (Disk: ctd=c7t0d0; valid=None; devpath=/pci@0,0/pci8086,2829@d/disk@0,0;

```

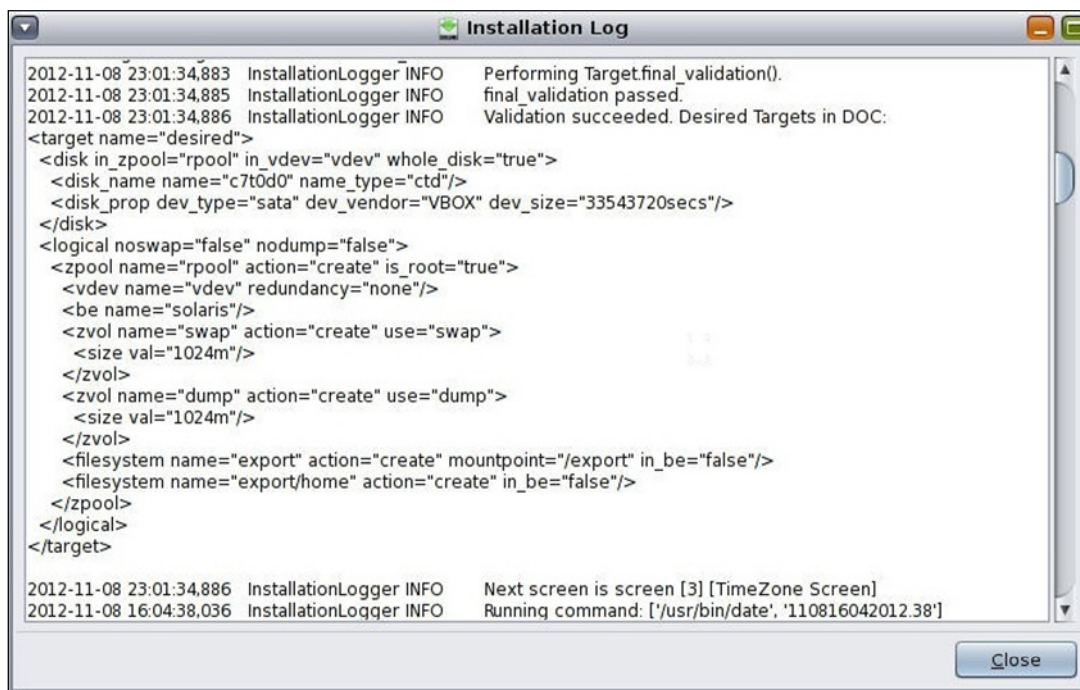
ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

The slide shows an example of an installation log. The log contains a complete record of each step of the installation process. Log files are an important tool in a system administrator's toolbox, so spend a few minutes acquainting yourself with the log content.

During the first part of the installation process, the configuration settings you supplied during the installation process are captured and applied to the target device.

Reviewing the Installation Log

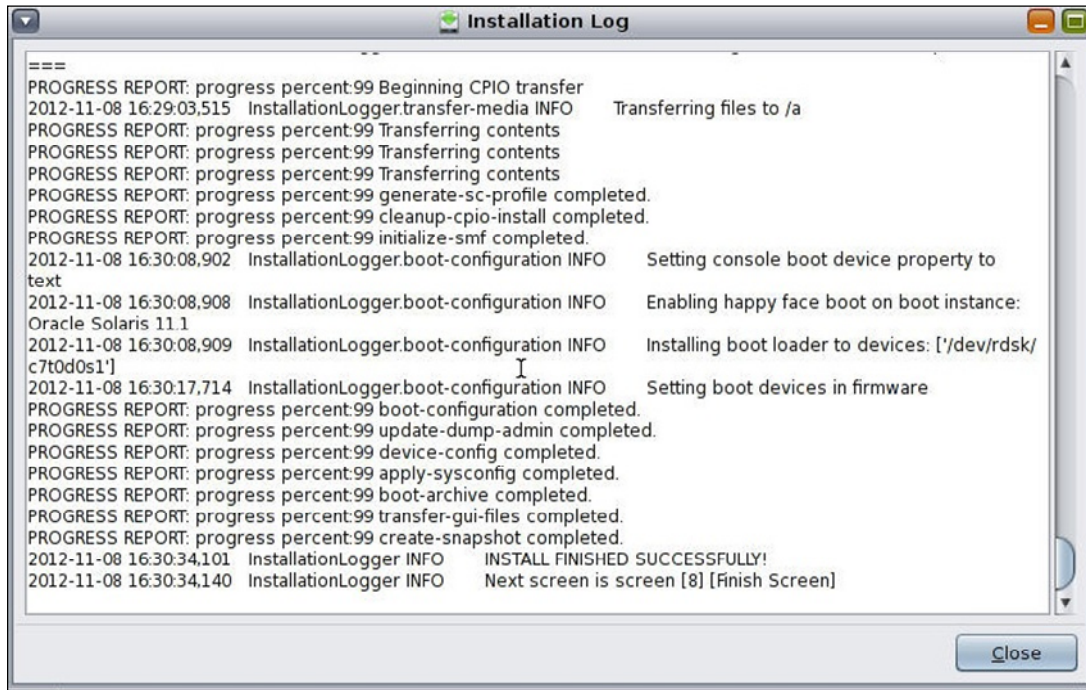


ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

This portion of the log shows the creation of the root pool (`rpool`).

Reviewing the Installation Log



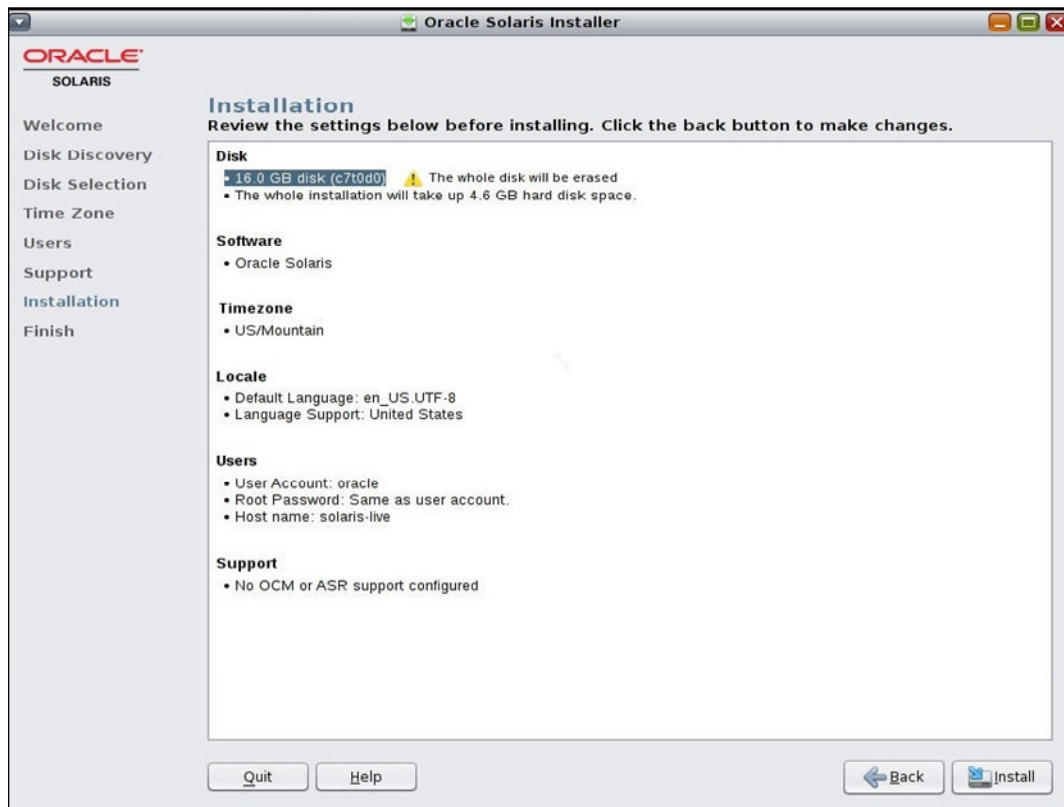
ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

During the last stage of the installation process, boot configuration takes place. The final step is the creation of a snapshot that captures the state of the system at this particular time. As you can see, the installation completed successfully.

After reviewing the installation log and verifying that no error messages were generated, you can return to the Finish screen by clicking **Close**.

Rebooting the System



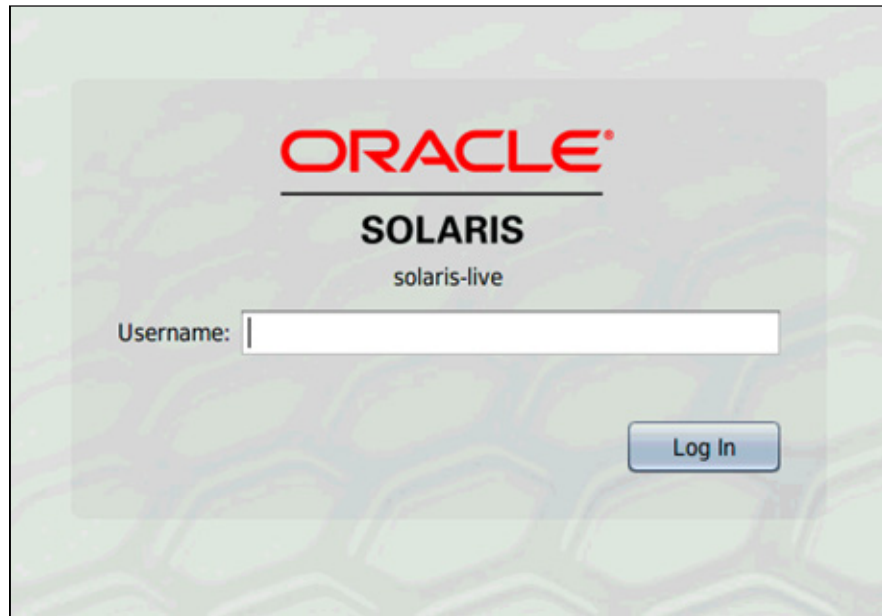
ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

After verifying that the installation was successful, you can reboot the system by clicking **Reboot**, or you can exit the installer and shut down the system.

Note: After the reboot, you can find the installation log at `/var/sadm/system/logs/install_log`.

Login Screen

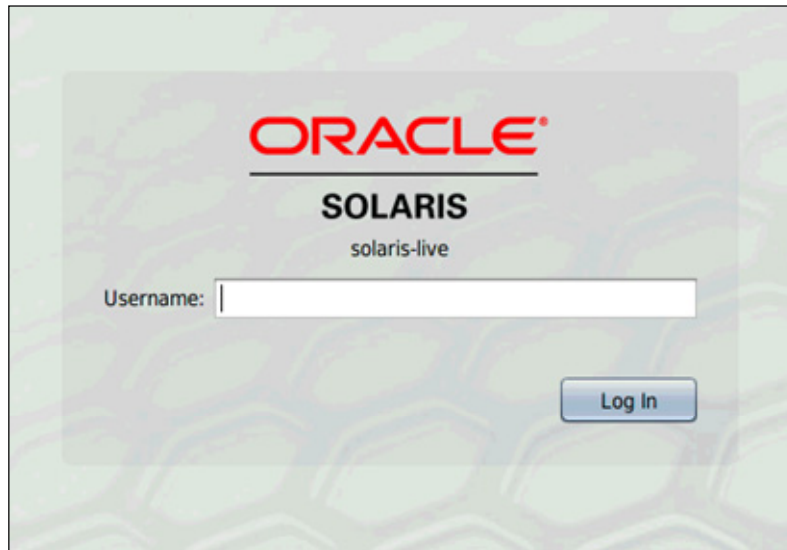


ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

After the system has rebooted, you should see the login screen.

Checking the Login Username

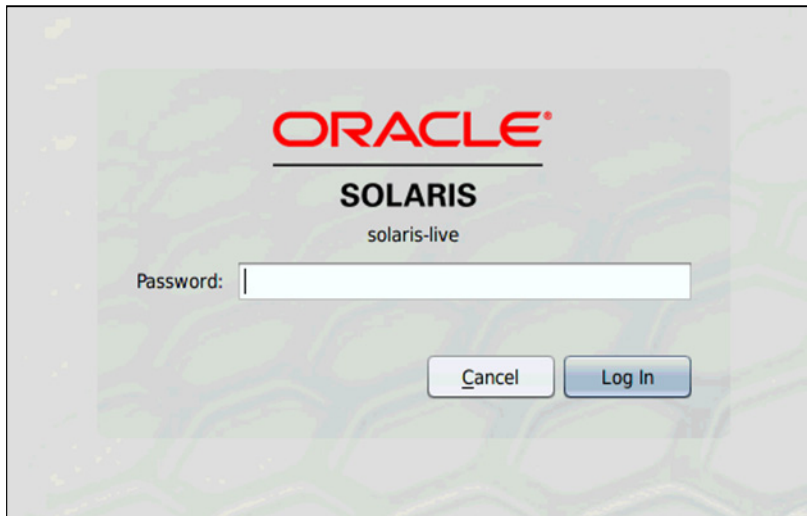


ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

On the login screen, enter the username that you created during installation and then click **Log In**. If the username is correct, the password screen appears.

Checking the Login Password



ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

To check the login password, enter the password that you provided during the installation. Then click **Log In**, as shown in the slide.

Lesson Agenda

- Performing pre-installation tasks prior to installing Oracle Solaris 11
- Installing Oracle Solaris 11 on a single host by using the Live Media Installer
- Installing Oracle Solaris 11 on a single host by using the Text Installer
- Installing Oracle Solaris 11 on multiple hosts by using the Automated Installer

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Installing Oracle Solaris 11 by Using the Text Installer

- You can perform an interactive text installation on individual SPARC and x86 client systems.
- If you have set up your network for automated installations, you can perform a text installation over the network.

Demo: Click [here](#) to view the procedure.

PDF: Click [here](#) to download a PDF copy of the procedure.

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Verifying the Installation

After installing Oracle Solaris 11 on the system, you should verify the installation. You should also gather key information about the system that can be used as a baseline for change-management documentation.

- Verifying login information
- Verifying the system's host name
- Displaying basic system information
- Displaying the system's release information
- Displaying disk configuration information
- Displaying the installed memory size
- Displaying information about network services
- Displaying network interface information

The Oracle logo, consisting of the word "ORACLE" in a white, sans-serif font, is positioned on the right side of a solid red horizontal bar.

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Verifying Login Information

- On the console login screen, enter the username that you established during installation. Then press Enter.
- If the username is correct, the password prompt appears.
- If you enter the correct password, you see the operating system release information and the command-line user prompt.

```
SunOS Release 5.11 Version 11.1 64-bit
Copyright (c) 1983, 2012, Oracle and/or its affiliates. All
rights reserved.
Loading smf(5) service descriptions: 199/199
Configuring devices.
Hostname: solaris-text

Solaris-text console login: solaris
Password:
Oracle Corporation SunOS 5.11 11.1 September 2012
oracle@solaris-text:~$
```



Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

At the console login (as shown in the code example), enter the username that you created during installation and then press **Enter**. If the username is correct, the password prompt appears.

To check the login password, go to the password prompt and enter the password that you provided during installation and press **Enter**. If the password is correct, you will see the operating system release information and the command-line user prompt (as shown in the example).

Verifying the System's Host Name

To display the host name, use `hostname`.

```
$ hostname  
solaris
```

The host name should match the computer name that you provided during installation.

The Oracle logo, consisting of the word "ORACLE" in a white, sans-serif font, is centered on a solid red rectangular background.

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

The name that appears here should match the computer name that you provided during installation. The host name is the name by which the system is known to a network. In this example, the host name is `solaris`.

Displaying Basic System Information

To display basic information about the system, use `uname -a`.

```
$ uname -a
SunOS solaris 5.11 11.1 i86pc i386 i86pc
```

In the output:

- **Operating system:** SunOS
- **Hostname:** solaris
- **Release:** 5.11
- **Version:** 11.1
- **Hardware name:** i86pc i386
- **Processor type:** i386

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Use the `uname -a` command to display the information listed in the slide. Capturing this information about your system for baseline purposes is extremely important for updating software packages.

Note: You can also run the `uname` command with specific options to display any one of the information items individually. See the `uname (1)` man page for details.

Displaying the System's Release Information

To display the operating system's release information, use `cat /etc/release`.

```
$ cat /etc/release
Oracle Solaris 11.1 X86
Copyright (c) 1983, 2012, Oracle and/or its affiliates.
All rights reserved.
Assembled 19 September 2012
```

The Oracle logo, consisting of the word "ORACLE" in a white, sans-serif font, centered on a red rectangular background.

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

The Oracle Solaris 11 operating system has a file called `/etc/release` that contains information about the system: the full operating system name, version of the release, hardware architecture, copyright, and date on which the release was assembled. You can use `cat /etc/release` to display the content of this file.

Note: Use the `cat /etc/release` command (instead of `uname -a`) to get current update release information.

Displaying Disk Configuration Information

To display disk information, switch to superuser and use `format`.

```
$ su -
Password:
# format
Searching for disks...done
AVAILABLE DISK SELECTIONS:
0. c7tod0 <ATA- VBOX HARDDISK-1.0-16.00GB>
    /pci@0,0/pci8086,2829@d/disk@0,0
Specify disk (enter its number):
selecting c7t0d0
[disk formatted]
/dev/dsk/c7t0d0s1 is part of active ZFS pool rpool. Please see
zpool(1M).

<continued on next page>
```

Note: The `format` utility requires superuser privileges.

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Use the `format` command to display information about the disks that are on the system. `format` is a utility for disk partitioning and maintenance. It enables you to format, label, repair, and analyze data, and to scrub data off a disk. To execute the `format` command, you must have root privileges. To switch to `root`, you can use the command `su -`.

Note: `su` stands for “switch user”.

In the example in the slide, only one disk has been configured in the system: `c7t0d0`. To see how a disk is formatted, enter the number of the disk (in this example, 0) after the `Specify disk (enter its number)` prompt, and then press **Enter**.

Displaying Disk Configuration Information: Format Menu

To display disk partition information, select `verify`.

```
FORMAT MENU:
  disk      - select a disk
  type      - select (define) a disk type
  partition - select (define) a partition table
  current   - describe the current disk
  format    - format and analyze the disk
  fdisk     - run the fdisk program
  repair    - repair a defective sector
  label     - write label to the disk
  analyze   - surface analysis
  defect    - defect list management
  backup    - search for backup labels
  verify    - read and display labels
  save      - save new disk/partition definitions
  inquiry   - show vendor, product and revision
  volname   - set 8-character volume name
  !<cmd>    - execute <cmd>, then return
  quit
format> verify
<continued on next page>
```

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

After you specify a disk, the **Format Menu** appears (as shown in the slide). To see how the disk has been partitioned (that is, how the slices on the disk have been allocated), enter **verify** at the `format` prompt, and then press **Enter**.

Displaying Disk Configuration Information: Partition Table

Primary label content:

Volume name = < >
 ascii name = <ATA-VBOX HARDDISK-1.0-16.00GB>
 bytes/sector = 512
 sectors = 33554431
 accessible sectors = 33554398

Part	Tag	Flag	FIRST Sector	Size	Last Sector
0	BIOS_boot	wm	256	256.00MB	524543
1	usr	wm	524544	15.74	33538014
2	backup	wu	0	0	0
3	unassigned	wm	0	0	0
4	unassigned	wm	0	0	0
5	unassigned	wm	0	0	0
6	unassigned	wm	0	0	0
7	unassigned	wm	0	0	0
8	reserved	wm	3353801	8.0MB	33554398

format>quit
 #

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

The partition table displays the partition-by-partition (or slice-by-slice) usage of disk space. The disk shown in this example has nine partitions, with the first partition (0) allocated to contain the root file system. Partitions 1, 3, 4, 5, 6, and 7 are currently unassigned. Partition 2 represents the whole disk, and partition 8 contains the boot program and is used in the booting process.

After displaying the current disk configuration, you can leave the `format` utility and return to the `root` user prompt by entering `quit` at the `format` prompt.

Displaying the Installed Memory Size

To display memory size, use `prtconf | grep Memory`.

```
# prtconf | grep Memory
Memory size: 1024 Megabytes
```

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

In Oracle Solaris, you can use the `prtconf` command to print the configuration that was initially created by the lower-level hardware discovery phase. This configuration includes the quantity of memory.

To display only the memory information, use `prtconf | grep Memory` (as shown in the example in the slide).

Because memory plays an important role in system performance, it is important to monitor the system's memory use.

Displaying Information About Network Services

To display information about network configuration services, use `svcs network/physical`.

```
# svcs network/physical
STATE      STIME      FMRI
online     15:35:09   svc:/network/physical:upgrade
online     15:35:09   svc:/network/physical:default
```

The Oracle logo, consisting of the word "ORACLE" in a white, sans-serif font on a red rectangular background.

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

The current network setup is important information to capture for your baseline documentation. To display network services information, use the `svcs network/physical` command.

In this example, observe that the `network/physical:default` service is online.

Displaying Network Interface Information

To display network interface information, use `ipadm show-addr`.

# ipadm show-addr			
ADDROBJ	TYPE	STATE	ADDR
lo0/v4	static	ok	127.0.0.1/8
net0/v4	dhcp	ok	10.0.2.15/24
lo0/v6	static	ok	::1/128
net0/v6	addrconf	ok	fe80::a00:27ff:fe4c:d1cb/10

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Network interface information is also important to capture for your baseline documentation. You want to make a note of the network interface information on your system (for example, `net0/v4` and `net0/v6`) and verify that these network interfaces are up and running.

To display the network interface information, use the `ipadm show-addr` command. The `ipadm` command is the tool for all IP interface configuration administration tasks.

Baseline System Information Commands: Summary

System Information	Command
Host name	hostname
Basic information: Operating system name, release, version, host name, hardware architecture, and processor type	uname -a
Operating system release information	cat /etc/release
Disk configuration	format
Installed memory	prtconf grep Memory
Information about network services	svcs network/physical
Network interface information	ipadm show-addr

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

The table in the slide summarizes the commands that display system information that can be used to document baseline information for your system. You can find out more about each of these commands in their respective man pages.

Quiz

You use the `cat /etc/release` command instead of `uname -a` to get current update release information.

- a. True
- b. False

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Answer: a

Lesson Agenda

- Performing pre-installation tasks prior to installing Oracle Solaris 11
- Installing Oracle Solaris 11 on a single host by using the Live Media Installer
- Installing Oracle Solaris 11 on a single host by using the Text Installer
- Installing Oracle Solaris 11 on multiple hosts by using the Automated Installer

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Automated Installer (AI): Overview

- The AI automates the installation of the OS on one or more SPARC and x86 systems over a network.
- The AI provides the following benefits:
 - Flexible configuration for disk layout, users, zone provisioning, and software selection
 - Support for unattended installation on multiple machines
 - Significant savings in installation time

The Oracle logo, consisting of the word "ORACLE" in white, uppercase letters on a red rectangular background.

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

The benefit of using the Automated Installer is that you can install and configure the operating system ISO on one server (either x86 or SPARC) and not have to replicate the same installation effort on other hosts.

When you use the AI, the operating system can be installed on client hosts unattended and without any manual intervention. This method saves significant installation time and is therefore used widely in the industry.

Automated Installer: Components


AI installation requires the following components:

- **IPS repository:** Provides the software packages that are identified in the AI manifest file to the client system
- **AI server:** Provides the install service that contains the installation instructions for the client system
- **DHCP server:** Provides the initial IP addresses and boot information
- **Client system:** Accesses the IP address information from the DHCP server and proceeds with automated installation

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Automated Installer: Process



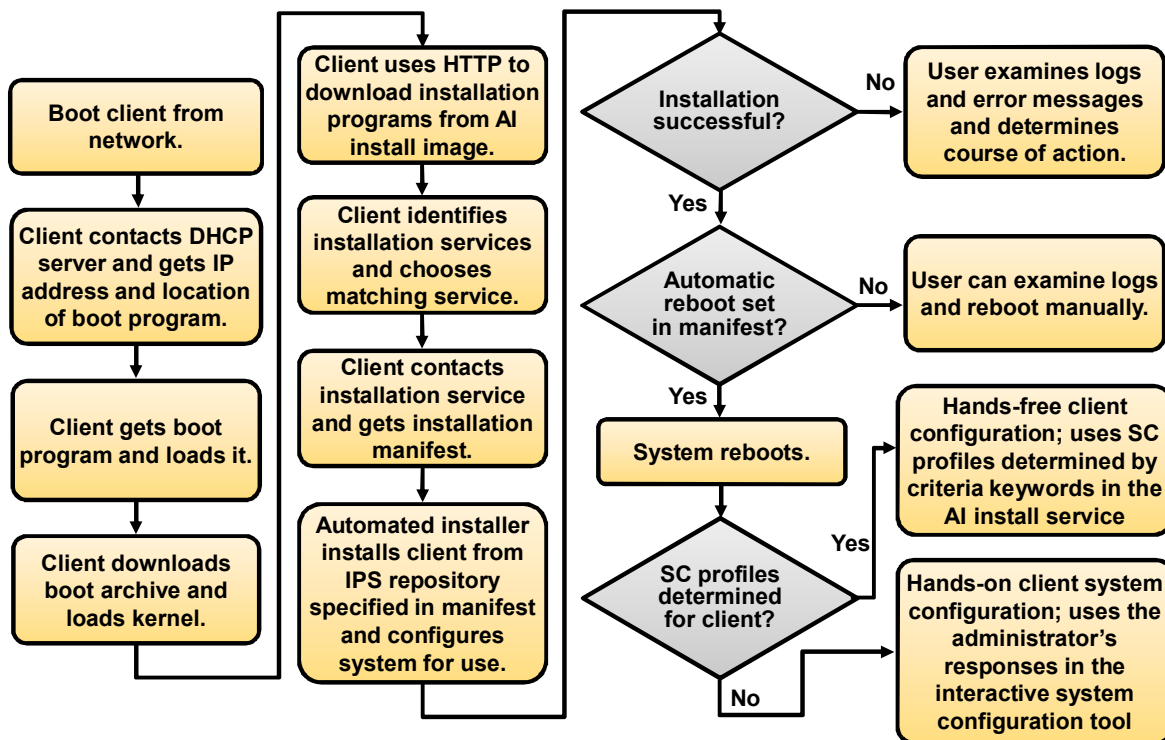
ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

The graphic in the slide illustrates the automated installation process:

1. A client system boots and receives IP and boot information from the DHCP server.
2. The client contacts an install service on the AI server and accesses the boot image and the AI manifest that contains the installation instructions.
3. The client is installed with the OS, pulling packages from the IPS repository that is specified in the installation instructions.

Automated Installer: Flowchart



ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

The flowchart in the slide illustrates the installation of a client system.

The client browses for available installation services, seeking a service where the installation criteria in the service's manifest file match the characteristics of the client system. When a match is found, the installation is performed on the client system by using a boot image and the manifest and system configuration (SC) profile specifications provided by the install service.

Performing an AI Installation

Performing an AI installation includes the following activities:

- Reviewing the AI installation server requirements
- Verifying the server software requirements
- Configuring the AI installation server
- Configuring the client system for an automated installation

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Reviewing AI Installation Server Requirements

Hardware	Requirement
Disk space	Approximately 0.75 GB additional disk space for each AI installation service after Oracle Solaris 11 OS has been installed
Memory	Recommended minimum: 1 GB

Software	Requirement
Operating system	Oracle Solaris 11 must be installed.
IP address	A static IP address must be used.
Router	The default route must be set.
DHCP	DHCP must be set up.
IPS repository	The repository must be set up locally.

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Before you start to configure a server as your AI server, you must check whether the server meets the following minimum requirements for an AI server.

Hardware Requirements

You must allow additional disk space for each AI installation service that you need. The disk space requirement listed in the slide is in addition to the disk space that you need for the Oracle Solaris 11 OS. The minimum requirement to operate an AI installation server is 1 GB of memory.

Software Requirements

You must have Oracle Solaris 11 installed. You must also configure the AI server to use a static IP address. If the AI server currently uses reactive network configurations, change the network configuration of the server for static IP addressing. AI clients rely on DHCP to obtain their initial IP addresses and boot files. You can configure the AI server to be the DHCP server by using the `installadm` command, or you can use a DHCP server that is already set up in the network.

A local IPS repository must be properly configured on your AI server to install the Oracle Solaris 11 OS on multiple network clients.

Verifying the Server Software Requirements

Check the following to ascertain that the server is ready to be configured as an AI server:

- Static IP address configuration
- Operational DNS
- DNS enabled for multicast
- IPS configured and available from the AI server
- Enabled DHCP server

The Oracle logo, consisting of the word "ORACLE" in a white, sans-serif font, is centered on a solid red rectangular background.

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

After selecting the server and installing the Oracle Solaris 11 OS, you must verify that the server meets the remaining software requirements. You do this by performing a series of checks. Ensure that the system has been configured with a static IP address and that domain name server (DNS) is operational. You must also ensure that an IPS repository and an IPS service are properly configured and available locally from the AI server. Finally, you ensure that the DHCP server is enabled.

Verifying the Static IP Address

To verify that the OS is configured with a static IP address, use `svcs network/physical` followed by `ipadm show-addr`.

# <code>svcs network/physical:default</code>			
STATE	STIME	FMRI	
online	15:02:57	svc:/network/physical:default	
# <code>ipadm show-addr</code>			
ADDROBJ	TYPE	STATE	ADDR
...			
net0/v4	static	ok	192.168.0.112/24
...			
#			

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

To verify that the OS is configured with a static IP address (that is, an IP address that has been created manually and not dynamically through the reactive network configuration or DHCP, for example), you first run the `svcs network/physical:default` command to verify that the physical network connection has been manually configured. In the example in the slide, the network has been set up manually and is online.

Next, you use the `ipadm show-addr` command to see the IP address information. In the code, there is a static network connection for the IP address 192.168.0.112/24, and the state of the connection is `ok`.

Verifying That DNS Is Operational

To verify that the DNS is operational, use `nslookup server domain name`.

```
# nslookup server1.mydomain.com
Server:          192.168.0.100
Address:         192.168.0.100#53

Name:   server1.mydomain.com
Address: 192.168.0.100
```

The Oracle logo, consisting of the word "ORACLE" in a white, sans-serif font, centered on a red rectangular background.

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

To verify that DNS is operational, use the `nslookup` command followed by the server's domain name.

Enabling the DNS Multicast Service

- To configure the AI server, ensure that the DNS multicast service is enabled.
- To enable the DNS multicast service, use `svcadm enable svc:/network/dns/multicast`.

```
# svcadm enable svc:/network/dns/multicast
```

- To verify that the service is enabled, use the `svcs network/dns/multicast` command.

```
# svcs network/dns/multicast
STATE          STIME          FMRI
online         1:32:27      svc:/network/dns/multicast:default
```



Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

DNS multicast enables the AI server to broadcast its IP–host name resolution information to the clients on the network.

Verifying That IPS Is Available Locally

- To verify that the correct local IPS repository is available to your server, use `pkg publisher`.

```
# pkg publisher
PUBLISHER    TYPE        STATUS      URI
solaris      origin      online      http://server1.mydomain.com
```

- To test IPS on the local server, search for the entire package by using `pkg search entire`.

```
# pkg search entire
INDEX        ACTION      VALUE        PACKAGE
pkg.fmri     set         solaris/entire  pkg:/entire@0.5.11-0.175.1.0.0.24.2
```



Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

For AI to work properly, it needs the IPS to be configured correctly and to be available from the AI server. To verify that a local IPS repository is available to the AI server, use the `pkg publisher` command and verify its URI.

To test that the IPS service is available, search for a given package by using the `pkg search` command.

Verifying That the DHCP Server Is Enabled

To verify that the DHCP server is enabled, use `svcs -a | grep dhcp`.

# <code>svcs -a grep dhcp</code>		
disabled	0:37:40	svc:/network/dhcp/relay:ipv6
disabled	0:37:40	svc:/network/dhcp/server:ipv6
disabled	0:37:40	svc:/network/dhcp/relay:ipv4
online	1:05:06	svc:/network/dhcp/server:ipv4

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Finally, ensure that the DHCP server is up and running. To do this, you use the `svcs -a | grep dhcp` command. The DHCP server should be in the `online` state.

Configuring the AI Server

The AI server is used to store an AI install image and contains the AI install service. Configuring an AI server involves the following steps:

1. Verifying the `netmasks` file configuration
2. Installing the AI installation tools
3. Setting up the AI boot image
4. Configuring an AI install service
5. Adding a client to the AI install service
6. Creating a custom manifest file
7. Configuring a custom SC profile

Demo: Click [here](#) to view the procedure.

PDF: Click [here](#) to download a PDF copy of the procedure.

The Oracle logo, consisting of the word "ORACLE" in a white, sans-serif font, is positioned on the right side of a solid red horizontal bar.

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

After verifying that the server meets the AI software requirements, proceed with configuring the AI. The AI server is used to store an AI install image and contains the AI install service. The AI install service specifies the installation instructions for installing the Oracle Solaris 11 OS on a client.

Installing the OS on the Client System

After configuring the AI server, proceed with an automated installation of the client system, which involves the following activities:

1. Identifying the client system requirements
2. Monitoring the installation remotely
3. Reviewing client installation messages

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Identifying Client System Requirements

Any system that meets the following requirements can be used as an AI client, including laptops, desktops, virtual machines, and enterprise servers.

Client System	Requirement
Disk space	Recommended minimum: 13 GB
Memory	Recommended minimum: 1 GB
Architectures	<ul style="list-style-type: none"> • X86: 64-bit only • SPARC: Oracle Solaris M-Series and T-Series systems only
Network access	<ul style="list-style-type: none"> • DHCP server that provides network configuration information • AI server • IPS repository that contains the packages to be installed on the client system

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

For automated installation over the network, SPARC and x86 client systems must meet the requirements listed in the table in the slide.

Additional SPARC client system requirements:

- **Firmware:** The firmware on SPARC clients must be updated to include the current version of the OpenBoot PROM (OBP) that contains the latest WAN boot support.
- **WAN boot:** SPARC clients of the AI installation over the network must support WAN boot.

Note: The recommended minimums are subject to change with the final release of the software. Check the Release Notes for final disk space and memory recommendations.

Identifying the Installation Files

- For x86 client installations, the `menu.lst` file is located in:
 - `/etc/netboot/menu.lst.01MAC_address` if `installadm create-client` was used
 - `/etc/netboot/<service_name>/menu.lst` if `installadm create-client` was not used
- For SPARC client installations, `system.conf` and `wanboot.conf` files are in:
 - `/etc/netboot/<service_name>`
 - For the `default-sparc` service, symlinks to these files are in `/etc/netboot`.

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

For x86 client systems, the `menu.lst` configuration file is created in the `/etc/netboot/` directory with one of the following file name formats:

- **`/etc/netboot/menu.lst.01MAC_address`:** If you used the `installadm create-client` command, the file name is `/etc/netboot/menu.lst.01MAC_address`, where `MAC_address` is the MAC address that was specified in the `installadm create-client` command.
- **`/etc/netboot/<service_name>menu.lst`:** If you did not use the `installadm create-client` command, the file name is `/etc/netboot/<service_name>menu.lst`, where `service_name` is the install service name that was specified in the `installadm create-service` command.

For SPARC client systems, the `system.conf` and `wanboot.conf` files are located in `/etc/netboot/<service_name>` if you have created an install service by using the `installadm create-service` command. For the `default-sparc` service, symlinks to these files are in `/etc/netboot`.

Note: You can remotely monitor an installation that is in progress by using Secure Shell. You do this by setting the `livessh` option to `enable` in the installation configuration file. After enabling access, log in to the AI client by using `jack` as the username and `jack` as the password.

Performing the Installation

To start the automated installation, boot the client.

- To boot an x86 client from the network, select the Oracle Solaris 11 11/11 Text Installer and command line boot option from the GNU GRUB menu.
- To boot a SPARC client and start an installation, use the following command from the OBP prompt:
`OK boot net:dhcp - install`

Demo: Click [here](#) to view the automated installation of the client system.

The Oracle logo, consisting of the word "ORACLE" in a white, sans-serif font, is centered within a solid red rectangular bar.

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

To start the automated installation, simply boot the client. The boot instructions for the SPARC and x86 clients are listed in the slide. When prompted, provide the configuration information.

Note: If you select the second install option shown in the GRUB menu, AI installation starts automatically.

Reviewing Client Installation Messages

- If the client successfully boots and downloads the install files, you see the Automated Installation started message.

```
The progress of the Automated Installation will be output to the console
Detailed logging is in the logfile at /system/volatile/install_log
Press RETURN to get a login prompt at any time.
```

- After the installation has completed successfully, you see the Automated Installation finished successfully message.

```
Automated Installation finished successfully
The system can be rebooted now
Please refer to the /system/volatile/install_log file for details
After reboot it will be located at /var/sadm/system/logs/install_log
```

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Log in as the `root` user with the password `solaris`. Monitor the progress of the installation via the installation log by using `tail -f /system/volatile/install_log`.

Note: To stop the `tail -f` command, press `Ctrl + C`.

After the installation has completed successfully, you have the option of rebooting. The client does not automatically reboot after a successful installation. You do, however, have the option of setting up an automatic reboot in the AI manifest. To enable automatic reboot, you set the `auto_reboot` attribute of the `<ai_instance>` tag to `true`. To reboot manually, run the `init 6` command.

If a client installation fails, the following list describes examples of actions that you can take based on the kind of errors you see.

Note: For complete details, go to `/system/volatile/install_log` to view the installation logs.

- **Check the connection to the IPS repository:** If a client cannot resolve the name of the IPS repository during installation, an error is generated. For this type of error, see if the client can ping the repository. If you get no response, you may have a connectivity problem. If the ping comes back as not having recognized the host, you may have a DNS problem.
- **Verify that DNS is configured on your client:** You can do this by verifying that a non-empty `/etc/resolv.conf` file exists. If this file does not exist or is empty, use `/sbin/dhclient` to check that your DHCP server is providing the DNS server information to the client. If this command returns nothing, the DHCP server is not set up properly. You must contact your DHCP administrator to correct the problem.
- **Check client boot errors:** Networking boot errors occur on both SPARC and x86 systems for a number of reasons, such as timing out issues or boot load failures. For more information about the types of errors that can occur and the possible causes of these errors (as well as suggested solutions), see *Installing Oracle Solaris 11 Systems*.

Quiz

A client can be associated with multiple install services.

- a. True
- b. False

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Answer: b

Summary

In this lesson, you should have learned how to:

- Perform pre-installation tasks
- Install Oracle Solaris 11 on a single host by using the Live Media Installer
- Install Oracle Solaris 11 on a single host by using the Text Installer
- Install Oracle Solaris 11 on multiple hosts by using the Automated Installer

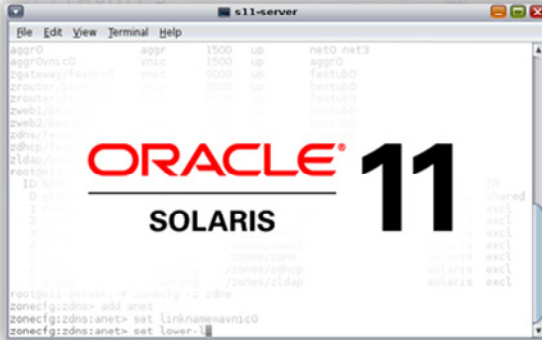
ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Monitoring System Resources

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.



System Administration for Experienced UNIX/Linux Administrators



**Administering System
Software by Using IPS**



**Administering Services
by Using SMF**



Administering ZFS



Configuring the Network



**Administering Oracle Solaris
Zones**



**Administering Privileges
and RBAC**



**Installing the Oracle Solaris 11
Operating System**



Monitoring System Resources

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Objectives

After completing this lesson, you should be able to:

- Identify the tools for monitoring system resources
- Explain the role of DTrace in diagnosing system issues

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Agenda

- Identifying the tools for monitoring system resources
- Explaining the role of DTrace in diagnosing system issues

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Monitoring and Observability Tools

The following generic utilities can be used to monitor and observe the system resources:

- I/O: `iostat`
- Kernel: `kstat`
- CPU: `mpstat`, `pgstat`
- Process: `prstat`, `truss`, `ptree`
- Virtual memory: `vmstat`
- SMF services: `svcs`
- File system: `fsstat`, `poolstat`
- Network: `netstat`, `dlstat`, `flowstat`, `ipmpstat`, `acctadm`
- Oracle Solaris Zones: `zonestat`

The Oracle logo, consisting of the word "ORACLE" in a white, sans-serif font, centered within a solid red rectangular background.

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

iostat Utility

- The `iostat` utility iteratively reports terminal, disk, and tape I/O activity, as well as CPU utilization.

```
# iostat [-cCdDeEiImMnpPrstxXYZ] [-l n] [-T u | d] /
[disk]... [interval [count]]
```

- For example, `iostat -xnp` generates partition and device statistics.

```
# iostat -xnp
extended device statistics
r/s  w/s  kr/s  kw/s  wait  actv  wsvc_t  asvc_t  %w  %b  device
0.4  0.3  10.4  7.9   0.0   0.0   0.0     36.9    0  1  c0t0d0
0.3  0.3  9.0   7.3   0.0   0.0   0.0     37.2    0  1  c0t0d0s0
0.0  0.0  0.1   0.5   0.0   0.0   0.0     34.0    0  0  c0t0d0s1
0.0  0.0  0.0   0.1   0.0   0.0   0.6     35.0    0  0  fuji:/export/ho
```

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

kstat Utility

- The `kstat` utility examines the available kernel statistics (kstats) about the system and displays the statistics that match the criteria specified on the command line.

```
# kstat [-lpq] [-T u | d ] [-c class] [-m module] [-i instance]
      [-n name] [-s statistic] [interval [count]]
```

- For example, the following command displays the CPU statistics of a system.

```
# kstat -p -m cpu_stat -s 'intr*'
cpu_stat:0:cpu_stat0:intr 4439717
cpu_stat:0:cpu_stat0:intrblk 858424
cpu_stat:0:cpu_stat0:intrthread 2216760
```

The Oracle logo, consisting of the word "ORACLE" in a white, sans-serif font, is centered on a solid red rectangular background.

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

mpstat Utility

- The `mpstat` utility reports per-processor or per-processor-set statistics.

```
# mpstat [-amq] [-A core|soc|bins] [-k keys] [-o num] [-p | -P
set] [-T d | u] [-I statfile] [-O statfile] [interval [count]]
```

- For example, `mpstat -ap 5` generates processor statistics over a five-second interval.

```
# mpstat -ap 5
SET minf mjf xcal intr ithr csw icsw migr smtx srw syscl usr sys wt idl sze
0      6   0  355  291  190  22   0   0   0   0   43   0   2   0  43   1
1     24  17  534  207  200  70   1   0   2   0  600   4   1   0  84   2
```

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

The command in the code example shows the processor set membership of each CPU. The default output is sorted by CPU number, aggregated by processor set, for user (`usr`) and system (`sys`) operations.

pgstat Utility

- The `pgstat` utility displays utilization statistics about processor groups (PGs).

```
# pgstat [-A] [-C] [-p] [-s key | -S key] [-t number] [-T u | d]
[-v] [-r string] [-R string] [-P pg ...] [-c processor_id... ]
[interval [count]]
```

- For example, `pgstat 1 2` displays utilization for all PGs over the last two seconds.

```
# pgstat 1 2
PG  RELATIONSHIP          HW      SW    CPUS
0   System                -    0.4%  0-31
3   Data_Pipe_to_memory   -    0.4%  0-31
2   Floating_Point_Unit   0%    0.4%  0-31
1   Integer_Pipeline      0%     0%   0-3
4   Integer_Pipeline      0%     0%   4-7
<output Truncated>...
```

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

A PG is a set of CPUs that are grouped together by a common characteristic. PGs are used by the operating system to represent CPUs that share performance-relevant hardware, such as execution pipelines and caches.

fsstat Utility

- The `fsstat` utility reports file system statistics.

```
# fsstat [-a|f|i|n|v] [-T | u|d] {-F | {fstype|path}...}  
[interval [count]]
```

- For example, `fsstat -f /` displays the statistics for each file operation for “/”.

```
# fsstat -f /  
/  
operation  #ops  bytes  
  open  8.54K  
  close  9.8K  
  read  43.6K  65.9M  
  write  1.57K  2.99M  
  ioctl  2.06K  
  setfl      4  
<Output Truncated>...
```

The Oracle logo, consisting of the word "ORACLE" in a bold, sans-serif font, with a registered trademark symbol (®) to the upper right.

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

poolstat Utility

- The `poolstat` utility reports active pool statistics.

```
# poolstat [-p pool-list] [-r rset-list] [interval [count]]
```

- For example, `poolstat -r pset` reports resource set statistics.

```
# poolstat -r pset
id pool          type rid rset          min  max size used load
0 pool_default  pset  -1 pset_default    1  65K   2  1.2  8.3
1 pool_admin    pset   1 pset_admin      1    1    1  0.4  5.2
2 pool_other    pset  -1 pset_default    1  65K   2  1.2  8.3
```

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

svcs Utility

- The `svcs` utility displays information about service instances as recorded in the service configuration repository.

```
# svcs [-aHpv?] [-o col[,col]]... [-R FMRI-instance]...  
[-sS col]... [FMRI | pattern]...
```

- For example, `svcs -x` displays service states.

```
# svcs -x  
svc:/application/print/server:default (LP print server)  
State: disabled since Mon Feb 13 17:56:21 2013  
Reason: Disabled by an administrator.  
See: http://support.oracle.com/msg/SMF-8000-05  
See: lpsched(1M)  
Impact: 2 dependent services are not running. (Use -v for list.)
```

The Oracle logo, consisting of the word "ORACLE" in a bold, sans-serif font, with a registered trademark symbol (®) to the upper right.

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

netstat Utility

- The `netstat` utility displays types of network data depending on the selected command-line option.

```
# netstat [-m] [-n] [-s] [-i | -r] [-faddress-family]
```

- For example, `netstat -i` displays the state of the interfaces.

```
# netstat -i
```

Name	Mtu	Net/Dest	Address	Ipkts	Ierrs	Opkts	Oerrs	Collis	Queue
lo0	8232	loopback	localhost	512	0	512	0	0	0
net0	1500	s11-client	s11-client	92	0	26	0	0	0
net1	1500	default	0.0.0.0	0	0	0	0	0	0
net3	1500	0.0.0.0	0.0.0.0	33	0	75	0	0	0

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

dlstat Utility

- The `dlstat` utility reports run time statistics about datalinks.

```
# dlstat [-r] [-t] [-Z] [-i interval] [-z zone[,...]] [link]
```

- For example, `dlstat` without any options displays the following statistics about the datalinks.

```
# dlstat
LINK      IPKTS      RBYTES      OPKTS      OBYTES
net1       0           0           0           0
net2       0           0           0           0
net0      198        62.93K       30         1.32K
net3       37         2.23K       181        61.90K
```

The Oracle logo, consisting of the word "ORACLE" in a white, sans-serif font on a red rectangular background.

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

flowstat Utility

- The `flowstat` utility reports runtime statistics about user-defined flows.

```
# flowstat [-r|-t] [-i interval] [-l link] [flow]
```

- For example, the following command displays traffic statistics for all flows at one-second intervals.

```
# flowstat -i 1
FLOW      IPKTS      RBYTES      IDROPS      OPKTS      OBYTES      ODROPS
udpflow   325.45K    39.5M       0           23.5K      7.3M       0
udpflow   123.37K    10.3M       0           11.7K      3.2M       0
udpflow    55.93K     3.9M        0           5.3K      1.4M       0
```

The Oracle logo, consisting of the word "ORACLE" in a bold, sans-serif font, with a registered trademark symbol (®) to its upper right.

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

ipmpstat Utility

- The `ipmpstat` utility displays information about the IPMP subsystem.

```
# ipmpstat [-n] [-o field[,...]] [-P]] -a|-g|-i|-p|-t
```

- For example, `ipmpstat -i` displays the following datalink statistics for the configured IPMP groups.

```
# ipmpstat -i
INTERFACE    ACTIVE  GROUP   FLAGS    LINK    PROBE    STATE
net2         yes    ipmp0   - - - - - up      disabled ok
net1         yes    ipmp0   - - mbM - - up      disabled
```

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

acctadm Utility

The extended accounting facility in Oracle Solaris 11 can be used to set up network accounting on the system:

1. View the status of the accounting types that can be enabled by the extended accounting facility.

```
# acctadm [process | task | flow | net]
```

2. Enable extended accounting for network traffic.

```
# acctadm -e extended -f filename net
```

3. Verify that extended network accounting is activated.

```
# acctadm net
```

4. Use the `dlstat` and `flowstat` commands to obtain historical statistics about the network.

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Network accounting involves capturing statistics about network traffic in a log file, which is useful for the purposes of tracking, provisioning, consolidation, and billing. You can later refer to the log file to obtain historical information about network use over a period of time. To set up network accounting, use the extended accounting facility's `acctadm` command.

After you set up network accounting, you can use the `dlstat` and `flowstat` commands to obtain historical statistics about network traffic. Ensure that you have configured flows on the system before you start using these commands.

zonestat Utility

- The zonestat utility reports on CPU, memory, networking, and resource utilization for the running zones.

```
# zonestat [-z zonelist] [-r reslist] [-n namelist] [-T u | d |
i] [-R reports] [-q] [-x] [-p [-P lines]] [-S cols] interval
[duration [report]]
```

- For example, zonestat -r summary 5 prints a report about zone utilization at five-second intervals.

```
# zonestat -r summary 5
Collecting data for first interval...
Interval: 1, Duration: 0:00:05
SUMMARYInterval: 3, Duration: 0:00:15
SUMMARY                Cpus/Online: 1/1    PhysMem: 1023M  VirtMem:
2047M

                ---CPU---  --PhysMem-- --VirtMem-- --PhysNet--
                ZONE  USED %PART  USED %USED  USED %USED  PBYTE %PUSE
                [total] 1.00 100%   658M 64.3%   839M 41.0%   1431 0.00%
                [system] 0.18 18.9%   373M 36.5%   521M 25.4%    -    -
<Output Truncated>...
```

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

The zonestat utility can be used from within a zone for a localized view of zone system resources, or it can be used from the global zone for a system-wide view of zone system resources. You can also identify resource bottlenecks or misbehaving applications with this utility.

vmstat Utility

- The `vmstat` utility reports virtual memory statistics about kernel thread, virtual memory, disk, trap, and CPU activity.

```
# vmstat [-cipqsS] [disks] [interval [count]]
```

- For example, `vmstat -5` displays a summary of the system activity every five seconds.

```
# vmstat 5
kthr    memory              page            disk          faults          cpu
r  b  w swap  free re  mf  pi  p  fr  de  sr  s0  s1  s2  s3   in   sy   cs  us  sy  id
0  0  0 11456 4120 1   41 19  1   3   0   2   0   4   0   0   48  112  130   4  14  82
0  0  1 10132 4280 0    4 44  0   0   0   0   0  23   0   0  211  230  144   3  35  62
<Output Truncated>...
```

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

prstat Utility

- The `prstat` utility reports active process statistics.

```
# prstat [-acJLmRtTv] [-C psrsetlist] [-j projlist]
[-k tasklist] [-n ntop[,nbottom]] [-p pidlist] [-P cpulist]
[-s key | -S key ] [-u euidlist] [-U uidlist] [-z zoneidlist]
[-Z] [interval [count]]
```

- For example, the following command reports the five most active superuser processes running on CPU1 and CPU2.

```
# prstat -u root -n 5 -P 1,2 1 1
```

PID	USERNAME	SWAP	RSS	STATE	PRI	NICE	TIME	CPU	PROCESS/LWP
306	root	3024K	1448K	sleep	58	0	0:00.00	0.3%	sendmail/1
102	root	1600K	592K	sleep	59	0	0:00.00	0.1%	in.rdisc/1
250	root	1000K	552K	sleep	58	0	0:00.00	0.0%	utmpd/1
288	root	1720K	1032K	sleep	58	0	0:00.00	0.0%	sac/1
1	root	744K	168K	sleep	58	0	0:00.00	0.0%	init/1
TOTAL:		25, load averages: 0.05, 0.08, 0.12							

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

truss Utility

- The `truss` utility traces system calls and signals.

```
# truss [-fcaeilddE] [- [tTvX] [!] syscall ,...] [- [sS] [!] signal ,...] [- [mM] [!] fault ,...] [- [rw] [!] fd ,...] [- [uU] [!] lib ,... : [:] [!] func ,...] [-o outfile] command | -p pid[/lwps]...
```

- For example, the following command shows only a trace of the `open`, `close`, `read`, and `write` system calls.

```
# truss -t open,close,read,write find . -print >find.out
```

The Oracle logo, consisting of the word "ORACLE" in a white, sans-serif font, is centered on a solid red rectangular background.

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

ptree Utility

- The `ptree` utility prints the process trees containing the specified PIDs or users, including child processes.

```
# ptree [-a] [-c] [-z zone] [pid | user]...
```

- For example, the following command prints the process tree (including children of process 0) for processes that match the command name `ssh`.

```
# ptree -a 'pgrep ssh'
1      /sbin/init
  100909 /usr/lib/ssh/sshd
    569150 /usr/lib/ssh/sshd
      569157 /usr/lib/ssh/sshd
        569159 -ksh
          569171 bash
            569173 /bin/ksh
              569193 bash
```



Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Quiz

Suppose that you need statistics about network traffic for the purposes of tracking, provisioning, consolidation, and billing. Which of the following tools would you use to prepare a comprehensive report?

- a. vmstat
- b. acctadm
- c. dlstat
- d. truss
- e. flowstat
- f. kstat

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Answer: b, c, e

Agenda

- Identifying the tools for monitoring system resources
- Explaining the role of DTrace in diagnosing system issues

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

DTrace: Overview

- A comprehensive, dynamic tracing facility that is built into the Oracle Solaris operating system
- An observability technology that helps you examine the behavior of user programs as well as the operating system in development and in production
- Features:
 - Examines the entire software stack
 - Determines the root cause of performance problems
 - Tracks the source of aberrant behavior

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

DTrace is supported on the following platforms:

- Oracle Solaris 10
- Oracle Solaris 11
- Oracle Linux
- Mac OS X

DTrace: Capabilities

- Analysis and observability
 - Provides a powerful new system and a process-centric framework for real-time analysis and observability
- Safety and comprehensive monitoring
 - Provides over 50,000 prebuilt data-monitoring points, inspection kernels, and user space levels
- Flexibility
 - Enables you to create custom programs to dynamically instrument the system
 - There is no need to instrument, stop, or restart your applications.
 - DTrace enables users to dynamically create as many monitoring points as required.

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

DTrace: Components

- Probes
- Providers
- Consumers
- D programming language

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Probes

- *Probes* are programmable sensors or points of instrumentation placed all over the Oracle Solaris system.
- DTrace provides thousands of probes.
- Each probe is associated with an action.
- When the probe fires, certain defined actions are executed.
- A four-tuple (`provider:module:function:name`) uniquely identifies every probe.

Example: `fbt:zfs:arc_read:entry`

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

In each four-tuple (`provider:module:function:name`):

- The probe is made available by a provider
- The probe identifies the module and function that it instruments
- The function is the name of the software function that contains the probe
- The probe has a name

Providers

- Libraries of probes are called *providers*.
- Providers make probes available to the DTrace framework.
- These libraries instrument a specific area of the system or a mode of tracing.
- New providers are added with every release of the operating system.

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Consumers

- *Consumers* are user mode programs that call in to the underlying DTrace framework.
- Four consumers are available in the current version of Oracle Solaris:
 - DTrace (1M)
 - lockstat (1M)
 - plockstat (1M)
 - intrstat (1M)

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

D Language

- Tracing programs, also referred to as *scripts*, are written in the D programming language.
- The language is a subset of C with added functions and variables that are specific to tracing.
- Additional capabilities of the D language include the following:
 - Supports ANSI C operators, strings, pointers, struct, and unions
 - Consists of expressions based on built-in variables: `pid`, `execname`, `timestamp`, and `curthread`
 - Performs arithmetic only on integers in D programs (floating-point arithmetic is not permitted in D)
 - Consists of CLI and scripting modes

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

D program clauses are written as single straight-line statement lists that trace an optional, fixed amount of data. The D language provides the ability to conditionally trace data and modify control flow by using logical expressions called *predicates* that can be used to prefix program clauses.

Although `if` statements are not supported, conditional statements can be written with the ternary condition operator (`?`).

```
expression1 ? expression2 : expression3
```

`expression2` and `expression3` can themselves contain the `?` operator, thus allowing nested conditional statements.

DTrace Toolkit

- The DTrace Toolkit is a collection of over 230 DTrace scripts and one-liners for performance observability and troubleshooting.
- The toolkit contains:
 - Scripts
 - Man pages
 - Sample documentation
 - Notes files
 - Tutorials
- To install the DTrace toolkit:
 1. `gunzip` and `tar xvf` the file.
 2. Run `./install`.

The Oracle logo, consisting of the word "ORACLE" in a white, sans-serif font, is centered within a solid red rectangular bar.

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

You can download the toolkit at the following location:

http://blogs.oracle.com/observatory/entry/DTrace_toolkit

For information about getting started, read the *Installation Guide*.

Setup

```
PATH=$PATH:/opt/DTT/Bin
```

```
MANPATH=$MANPATH:/opt/DTT/Man
```

DTrace Toolkit: Important Scripts

Script Folder	Function
Apps	For certain applications (such as Apache and NFS)
CPU	Measuring CPU activity
Disk	Analyzing I/O activity
Extra	For other categories
Kernel	For kernel activity
Locks	Analyzing locks
Mem	Analyzing memory and virtual memory
Net	Analyzing activity of network interfaces and the TCP/IP stack
Proc	Analyzing activity of a process
System	Measuring system-wide activity
User	Monitoring activity by UID
Zones	Monitoring activity by zone

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

The table in the slide summarizes the folders that contain some of the more important DTrace Toolkit scripts.

Before Using DTrace

Consider the following options before using DTrace:

- Perform a sanity check first.
 - `/var/adm/messages`
- Start with the “Big 5” stat tools.
 - `vmstat`, `mpstat`, `iostat`, `prstat`, and `netstat` for memory leaks
- Answer the following high-level questions for a quick initial diagnosis:
 - Is there a significant number of cache misses?
 - How much time is spent in user mode compared with system mode?
 - Is the system short on memory or other critical resources?
 - What are the system’s I/O characteristics?

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Because DTrace produces a lot of data that might overwhelm you, it is useful to start with other tools and diagnostic options and then use DTrace for precision diagnosis.

Launching DTrace

Before launching DTrace, consider the following facts:

- Only `root` is allowed to run DTrace by default.
- Non-`root` users need one or more of the following privileges to access DTrace:
 - `DTrace_kernel`
 - `DTrace_proc`
 - `DTrace_user`

```
% ppriv -l | grep Dtrace
DTrace_kernel DTrace_proc DTrace_user
```

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

- `DTrace_kernel` enables you to examine the kernel.
- `DTrace_proc` enables you to trace your processes.
- `DTrace_user` enables you to examine the syscalls and profile the processes.

If an unprivileged user tries to run DTrace, the following error message is displayed:

```
$ id uid=1001(user1) gid=1(other)
$ /usr/sbin/DTrace -n 'syscall::exece:return' DTrace:
failed to initialize DTrace: DTrace requires additional privileges
```

Caution: DTrace enables users to perform powerful actions that can modify the state of a program because it enables visibility into all aspects of the system, including user-level functions, system calls, kernel functions, and more. Therefore, you must apply caution while privileging users. Only authorized users should be DTrace-enabled on a production system.

DTrace: Example

Scenario: Your system processes are causing a massive number of system calls and subsequently having a negative effect on system performance. You want to investigate the number of system calls that the processes are making.

```
# dtrace -n 'syscall::read:entry @[execname] = count()}'
dtrace: description 'syscall::read:entry ' matched 1 probe
^C
gnome-settings-d          1
in.mpathd                 2
gnome-terminal            3
init                      4
hald                      72
hald-addon-acpi           72
Xorg                      248
#
```

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

The example in the slide provides a count of the processes that are calling the `read` function. The `syscall` provider enables the `entry` probe in the `read` function to monitor the processes that are invoking the `read` system calls. You can make the investigation more granular by using specific DTrace one-liners or scripts.

The DTrace one-liner in the slide can also be written in the form of a D language script.

```
# vi read.d
#!usr/sbin/dtrace
syscall::read:entry
(
    @[execname] = count();
}
# chmod +x read.d
# dtrace -s read.d
```

Note: Traditional performance monitoring tools such as `truss` are mostly process-centric. Unless the processes are identified, it is difficult to address a performance bottleneck. DTrace provides the option to identify systemic problems and help isolate specific issues.

Summary

In this lesson, you should have learned how to:

- Identify the tools for monitoring system resources
- Explain the role of DTrace in diagnosing system issues

ORACLE

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

