

Oracle® Communications Session Border Controller

Platform Preparation and Installation Guide



Releases S-CZ8.1.0, E-CZ8.1.0
F20250-02
December 2020

The Oracle logo, consisting of a solid red square with the word "ORACLE" in white, uppercase, sans-serif font centered within it.

ORACLE®

F20250-02

Copyright © 2014, 2020, Oracle and/or its affiliates.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs (including any operating system, integrated software, any programs embedded, installed or activated on delivered hardware, and modifications of such programs) and Oracle computer documentation or other Oracle data delivered to or accessed by U.S. Government end users are "commercial computer software" or "commercial computer software documentation" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, reproduction, duplication, release, display, disclosure, modification, preparation of derivative works, and/or adaptation of i) Oracle programs (including any operating system, integrated software, any programs embedded, installed or activated on delivered hardware, and modifications of such programs), ii) Oracle computer documentation and/or iii) other Oracle data, is subject to the rights and limitations specified in the license contained in the applicable contract. The terms governing the U.S. Government's use of Oracle cloud services are defined by the applicable contract for such services. No other rights are granted to the U.S. Government.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

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

Intel and Intel Inside are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Epyc, and the AMD logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and documentation may provide access to or information about content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services unless otherwise set forth in an applicable agreement between you and Oracle. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services, except as set forth in an applicable agreement between you and Oracle.

Contents

About This Guide

1 Introduction to Platform Preparation and Software Deployment

Acme Packet Engineered Platforms	1-1
Virtual Machines	1-1
Netra and COTS	1-2

2 Netra Server X5-2 Platform Preparation

Cable the Netra X5-2 for Communications	2-2
Available Connections	2-4
Cable the Local Console	2-5
Connect ILOM to the Netra X5 for Communications	2-6
Software Installation - Netra and Server-based Platforms	2-7
Known Issues	2-9

3 Oracle Server X7-2 Platform Preparation

Cable the Oracle X7-2	3-7
Available Connections	3-9
Cable the Local Console	3-10
Connect ILOM to the Oracle X7-2	3-11
Software Installation - Oracle X7-2 Platforms	3-12
Known Issues	3-14

4 Session Delivery Products as Virtual Machines

Hyper-threading and CPU Affinity	4-1
Host Hypervisor CPU Pinning	4-1
Configuration Overview	4-2
VLAN Support	4-3

Provisioning Entitlements	4-3
5	Virtual Machine Platforms
Create and Deploy on Oracle VM Manager	5-1
Create and Deploy on KVM	5-3
Create and Deploy on VMware®	5-5
Create and Deploy on Hyper-V®	5-8
Configure OpenStack Heat Template	5-12
6	Boot Management
Boot Loader Overview	6-1
Boot Parameters	6-2
Boot Parameter Definitions	6-3
Boot Flags	6-4
Changing Boot Parameters	6-4
Change Boot Parameters from the ACLI	6-5
Change Boot Parameters by Interrupting a Boot in Progress	6-6
7	Formatting the Disk Volume
40GB or Less Format Plan	7-1
40GB or More Default Format Plan	7-1
40GB or More Custom Format Plan	7-1
Formatting Procedure	7-2
8	Interface Considerations for VM Platforms
Software Adaptations for Physical Interfaces	8-1
COTS Interfaces	8-1
COTS Network Interfaces	8-2
The interface-mapping Branch	8-2
Working with the interface-mapping branch	8-3
interface-mapping	8-3
Serial Interfaces	8-5
9	Flash Drive Installation via Boot Media Creator
Creating a Build Image	9-1

10 Software Upgrade

Upgrade Checklist	10-1
Download the Software	10-1
Stand-alone Upgrade	10-2
Check /boot for free space	10-2
Upload the Stage 3 Boot Loader and System Image	10-2
Software Upgrade Procedure	10-3
HA Upgrade	10-4
HA Upgrade Procedure	10-5
HA Backout Procedure	10-7

A Physical Interfaces on Acme Packet Platforms

Acme Packet 1100 Physical Interfaces	A-2
Acme Packet 3900 Physical Interfaces	A-3
SNMP Hardware Reporting	A-4
Acme Packet 3900 MIBS Paths	A-5
Acme Packet 4600 Physical Interfaces	A-5
Acme Packet 6100 Physical Interfaces	A-6
Acme Packet 6300/6350 Physical Interfaces	A-7
Acme Packet 6300	A-8
Acme Packet 6350	A-9

List of Figures

2-1	Netra X5-2 for Communications Configuration A (4 Onboard 10 GigE Ports)	2-3
2-2	Netra X5-2 for Communications Configuration B (4 Onboard 10 GigE Ports & 1 Quad GigE NIC)	2-3
2-3	Netra 5-2 for Communications Configuration C (4 Onboard 10 GigE Ports & 2 Quad GigE NICs)	2-3
2-4	Connecting to USB, VGA and SER MGT (COM1) Ports	2-6
2-5	Connecting to ILOM Port	2-7
3-1	Selecting RAID Configuration	3-2
3-2	Begin RAID Configuration	3-3
3-3	Clear Any Existing RAID Configuration	3-3
3-4	RAID - Create Virtual Drive	3-4
3-5	Set Drive to RAID1	3-4
3-6	RAID - Select Drives	3-5
3-7	Select All Drives	3-5
3-8	Save RAID Configuration	3-6
3-9	Initialize RAID Configuration	3-6
3-10	Exit RAID Configuration	3-7
3-11	Oracle X7-2 Configuration A (4x10 GigE NIC)	3-8
3-12	Oracle X7-2 Configuration B (Two 4x10 GigE NICs)	3-8
3-13	Oracle X7-2 Configuration B (One QSFP and One 4x10 GigE NICs)	3-8
3-14	Connecting to USB and SER MGT (COM1) Ports	3-11
3-15	Connecting to ILOM over the Network	3-11
4-1	Contrast of Data Paths with and without CPU Pinning	4-2
5-1	Properties Dialog for OVF Deployment	5-7
5-2	OVF Deployment Status Dialog	5-7
5-3	VM Settings on Hyper-V	5-11
9-1	BMC - Step 3	9-2
9-2	BMC - Step 4	9-2
9-3	BMC - Step 5	9-3
9-4	BMC - Step 6	9-3
9-5	BMC - Step 7	9-4
9-6	BMC - Step 8	9-4
9-7	BMC - Step 9	9-5
9-8	BMC - Step 10	9-6
10-1	Configure and Reboot Standby First for HA Procedures	10-5

A-1	Acme Packet 1100 - Rear View	A-2
A-2	Acme Packet 3900 - Rear View	A-3
A-3	Acme Packet 4600 - Rear View	A-5
A-4	Acme Packet 6100 - Rear View	A-7
A-5	Acme Packet 6300/6350 Management Interfaces	A-8
A-6	Acme Packet 6300 Media Interfaces	A-9
A-7	Acme Packet 6350 - Media Interface	A-10

About This Guide

This Installation and Platform Preparation Guide addresses platform preparation of platforms that support Oracle Communications Session Delivery Product version CZ8.1.0. This software encompasses multiple products. This content bridges the information between physical installation and initial power on procedures. The guide takes into account the fact that a variety of platforms are supported, each one posing its own operational considerations, and presents information specific to each platform that enables proper product operation.

Whereas platform documentation for physical installation is provided by Oracle and the respective vendors, this guide addresses those details in between physical installation and service configuration. For service configuration, see the applicable Configuration Guide for your product. In addition, the information herein can help users after product deployment to, for example, identify physical interfaces.

Related Documentation

The following table describes the documentation set for this release.

Document Name	Document Description
Acme Packet 4600 Hardware Installation Guide	Contains information about the components and installation of the Acme Packet 4600.
Acme Packet 6100 Hardware Installation Guide	Contains information about the components and installation of the Acme Packet 6100.
Acme Packet 6300 Hardware Installation Guide	Contains information about the components and installation of the Acme Packet 6300.
Acme Packet 6350 Hardware Installation Guide	Contains information about the components and installation of the Acme Packet 6350.
Release Notes	Contains information about the current documentation set release, including new features and management changes.
ACLI Configuration Guide	Contains information about the administration and software configuration of the Service Provider Oracle Communications Session Delivery Product.
ACLI Reference Guide	Contains explanations of how to use the ACLI, as an alphabetical listings and descriptions of all ACLI commands and configuration parameters.
Maintenance and Troubleshooting Guide	Contains information about Oracle Communications Session Delivery Product logs, performance announcements, system management, inventory management, upgrades, working with configurations, and managing backups and archives.

Document Name	Document Description
MIB Reference Guide	Contains information about Management Information Base (MIBs), Oracle Communication's enterprise MIBs, general trap information, including specific details about standard traps and enterprise traps, Simple Network Management Protocol (SNMP) GET query information (including standard and enterprise SNMP GET query names, object identifier names and numbers, and descriptions), examples of scalar and table objects.
Accounting Guide	Contains information about the Oracle Communications Session Delivery Product's accounting support, including details about RADIUS and Diameter accounting.
HDR Resource Guide	Contains information about the Oracle Communications Session Delivery Product's Historical Data Recording (HDR) feature. This guide includes HDR configuration and system-wide statistical information.
Administrative Security Essentials	Contains information about the Oracle Communications Session Delivery Product's support for its Administrative Security license.
Security Guide	Contains information about security considerations and best practices from a network and application security perspective for the Oracle Communications Session Delivery Product family of products.
Installation and Platform Preparation Guide	Contains information about upgrading system images and any pre-boot system provisioning.
Call Traffic Monitoring Guide	Contains information about traffic monitoring and packet traces as collected on the system. This guide also includes WebGUI configuration used for the SIP Monitor and Trace application.
Header Manipulation Rule Guide	Contains information about configuring and using Header Manipulation Rules to manage service traffic.

Revision History

Date	Description
April 2018	<ul style="list-style-type: none"> Initial release
May 2018	<ul style="list-style-type: none"> Removes extraneous content on 4500 disk format after downgrade
October 2019	<ul style="list-style-type: none"> Updates port numbers in the table in "Acme Packet 4600 Physical Interfaces."
February 2020	<ul style="list-style-type: none"> Updates the Media Interface port numbers in "Acme Packet 1100 Physical Interfaces."
December 2020	<ul style="list-style-type: none"> Updates INT and EXT port speeds in "Acme Packet 1100 Physical Interfaces."

1

Introduction to Platform Preparation and Software Deployment

This documentation explains platform preparation, power-on and software deployment to the point where the user is ready to perform service configuration. Refer to Hardware Installation documentation for rack and stack procedures.

Platform support by Oracle Communications session delivery software extends to generic platforms. As such, preparation of these platforms requires that the user perform tasks independent of their session delivery product. Although the user needs to use platform documentation for this platform-specific information, Oracle distributes this documentation to provide the user with:

- Settings required by the session delivery software.
- Guidance to procedures that apply to session delivery software.

Acme Packet Engineered Platforms

Acme Packet engineered platforms are shipped with software pre-loaded on the system. When you power up the hardware for the first time, it will boot with a product image ready for configuration.

To install a newer patch release, or different version of software all together, continue to the Boot Management and Software Upgrade chapters in this guide for all considerations and procedures required to a system update.

Virtual Machines

Virtual Machines (VMs) supported by Oracle Communications Session Delivery software varies across software version. Find specific version support within the context of your version's documentation.

Operation over VMs is roughly equivalent to deployment over COTS/Server hardware. Platform preparation, however, differs greatly. In addition, platform preparation differs greatly between VM platforms.

Preparation procedures that apply to all VM platforms include the following steps:

1. Make the VM template available to the VM manager.
2. Configure the VM manager to apply the template correctly for Oracle Communications Session Delivery software.
3. Power-on the VM. If the deployment is using a VM template, the system uses that template to automatically install onto the virtual drive, after which the server reboots. Deployments using raw images do not perform an installation process.

VM deployment requires extensive knowledge about the specific platform that is not documented herein. The intent of this documentation is to provide information that

helps the user navigate the deployment and perform tasks that are specifically related to Oracle Communications Session Delivery software.

Netra and COTS

You must provision Netra (and COTS) hardware before installing software. This includes platform-specific configuration, such as BIOS, and platform management access. Once a version of the Oracle Communications Session Delivery Product is running on a Netra or COTS platform, you simply install new software using the instructions provided in the Boot Management and Software Upgrade chapters herein.

2

Netra Server X5-2 Platform Preparation

Oracle Communications produces a variety of software products that run on the Netra X5-2 for Communications platform, including Oracle session delivery applications.

Use your Hardware documentation to install and establish system management by way of ILOM. Then use the steps below to prepare the Netra X5-2 for session delivery software installation.

Note:

The [Connect ILOM to the Netra X5 for Communications](#) also displays ILOM cabling.

1. Confirm applicable firmware on the server.
 - To check the firmware versions installed in the server, go to the Oracle Integrated Lights Out Manager (ILOM) web interface, and navigate to **System Information, Firmware**.
 - Review your session delivery product Release Notes for qualified software and firmware versions.
2. Upgrade or downgrade the firmware on the server as necessary. Go to https://docs.oracle.com/cd/E37444_01/index.html for ILOM upgrade instructions.
3. Configure the BIOS settings. (Settings navigation may differ based on the BIOS version.)
 - a. Observe the boot procedure and use the documented key sequence to interrupt the boot and display the BIOS configuration dialogs. For example, pressing the F2 key is a common way to enter BIOS configuration from a terminal application that supports function keys.
 - b. Navigate to the Boot menu and, depending on the software distribution you are using, set the USB or CD as the first device followed by the disk controller.
 - c. Disable Hyper-Threading.
 - d. Disable CPU power limit.
 - e. Disable C6 Reporting.
 - f. Disable the UEFI Stack.
 - g. Change Energy Performance to Performance. (For example, set "ENERY_PERF_BIAS_CFG" mode to "PERF".)
 - h. To decrease boot up time, Oracle recommends disabling Intel PXE Boot Agent for both onboard and NIC ethernet ports. To disable the Boot Agent for the onboard ethernet ports, navigate to the OpROM option for NET0, NET1, NET2, and NET3 interfaces (for example, IO, Internal Devices) and set it to disabled.

- i. To disable Boot Agent for NIC ethernet ports, note the blue PCIe slot number label at the back of the Netra server where the NICs are installed, then disable the OpROM option for those slots. (Note that you may be able to identify slot number through the ILOM System Information, PCI Devices menu.)
 - j. Reboot the server.
4. Initialize the HDD.
 - a. Open the ILOM remote system console to observe the system's boot cycle, and interrupt the boot cycle to enter the LSI MegaRAID status display. For example, pressing the Ctrl-R key is a common way to enter LSI MegaRAID BIOS Configuration Utility.
 - b. Navigate the utility to establish the elements of your virtual drive, typically consisting of a New Configuration with two entire HDDs.
 - c. Access the menu from which you create a virtual drive.
 - d. Set the RAID level to RAID-1.
 - e. Select all of the drives that you want.
 - f. From the Virtual Drive Management dialog, select the new drive and initialize it. For example, pressing F2 and selecting Fast Init from the command menu is a common way to execute initialization.
 - g. After initialization is complete, Escape from the LSI MegaRAID Configuration Utility and reboot the system.
5. Perform a cold shutdown by removing all system power.

Cable the Netra X5-2 for Communications

After mounting the Netra X5-2 for Communications in an equipment rack and installing all components, use the following instructions to connect all appropriate data cables to the ports before powering the system up and beginning the configuration.

Oracle has qualified the following configurations of the Netra X5-2 for Communications (the onboard 10 GigE ports are configured for 1G operation):

- Configuration A: Four onboard 10 GigE ports and no Quad GigE NIC
- Configuration B: Four onboard 10 GigE ports and 1 Quad GigE NIC
- Configuration C: Four onboard 10 GigE ports and 2 Quad GigE NICs

Figure 2-1 Netra X5-2 for Communications Configuration A (4 Onboard 10 GigE Ports)




Figure 2-2 Netra X5-2 for Communications Configuration B (4 Onboard 10 GigE Ports & 1 Quad GigE NIC)





Figure 2-3 Netra 5-2 for Communications Configuration C (4 Onboard 10 GigE Ports & 2 Quad GigE NICs)



Oracle recommends using Category 6 (or better) for all Ethernet connections.

You do not need to use every port for proper operation.

You can install and remove Ethernet and 1000BASE-T cables while the Netra X5-2 for Communication runs, but when you disconnect a cable the link is lost and the system generates an alarm.

Available Connections

Please read all of the information for each of the available connections prior to cabling the Netra X5-2 for Communications.

Port	Description	You Need
NET (0-3)	10 GigE ports - labeled Net 3, Net 2, Net 1 and Net 0 (left to right). Enables you to connect the Netra X5-2 to your network.	A Category 6 (or better) Ethernet cable to connect to the NET 0 port to your network Network parameters such as an IP address (can be provided by DHCP services or assigned a static address in the OS) Additional Category 6 (or better) Ethernet cables and Ethernet addresses as needed for additional connections to NET 1 - 3
NET MGT	Provides a 10/100BASE-T Ethernet connection to the Service Processor (SP) through an RJ-45 connector. The NET MGT port provides support connections to the SP using the Oracle Integrated Lights Out Manager (ILOM) CLI and Web interface. By default, the NET MGT port is configured to use DHCP to automatically obtain an IP address. Alternatively, you can assign a static IP address to the NET MGT port. To use the NET MGT port, you must configure its network settings. Once configured, use the NET MGT port IP address to log on to the device using a browser or secure shell.	Category 6 (or better) Ethernet cable to connect the NET MGT port to your network IP address for this port (required from DHCP or a static address)

Port	Description	You Need
SER MGT (COM1)	Provides a TIA/EIA-232 serial Oracle/Cisco standard connection to the SP through an RJ-45 connector. Default settings: 8N1: eight data bits, no parity, one stop bit 115200 baud Disable hardware flow control (CTS/RTS) Disable software flow control (XON/XOFF)	A terminal device (e.g., terminal, connection to a terminal server, or computer such as a laptop running terminal emulation software) A cable to connect the terminal device to the SER MGT (COM1) port
USB	Provides USB connections to the SP. The USB ports are hot pluggable, so you can connect and disconnect USB cables from these ports and peripheral devices without affecting server operations.	USB keyboard USB mouse Note: Maximum USB cable length: 5 meters
VIDEO	Provides a temporary video connection to the SP.	VGA monitor HDB-15 video cable with a maximum cable length of 6 meters (19.7 feet)

Cable the Local Console

You can connect the Administration console to either the Oracle Integrated Lights Out Manager (ILOM) (NET MGT), the local VGA+USB console ports, or the local SER MGT (COM1) serial console port.

- To cable a serial console connection:
 - Serial console cable with an RJ-45 connector
- To cable a USB and Video Connection:
 - DB-15 video cable with a maximum cable length of 6 meters (19.7 feet)
 - USB cable with a maximum cable length of 6 meters (19.7 feet)
 - USB keyboard

In the following procedure, you have the option to either cable a serial connection or to cable a USB/Video connection.


 **Note:**

Do not configure COM2 in the bootparams menu.

When configuring boot loader parameters, set the console to VGA when you use ILOM or VGA+USB, or to COM1 when you use SER MGT. The boot loader is accessible on all console ports, but only input from the active console port can be recognized by the Netra X5-2 for Communications.

1. Locate the appropriate cables to connect to the Netra X5-2 for Communications.
2. To cable a serial connection, insert the serial console cable into the SER MGT (COM1) port.

Figure 2-4 Connecting to USB, VGA and SER MGT (COM1) Ports



 **Note:**

Refer to the Netra X5-2 for Communications hardware documentation for information on how to configure the terminal application to connect to the console, and how to establish communications with the Netra X5-2 for Communications.

3. To cable a USB/Video connection, do the following:
 - a. Insert the 15-pin connector end of the video cable into the Video port.
 - b. Insert the USB cable from the mouse and keyboard into the USB ports.
4. Lead the cables neatly away from the rear panel.
5. Plug in the cables to their respective destination components.

Connect ILOM to the Netra X5 for Communications

Use the following procedure to make a connection to the Netra X5-2 for Communications Oracle Integrated Lights Out Manager (ILOM) port. For a remote permanent connection to the Service Processor over the ILOM connection, use the rear panel NET MGT port.


 **Note:**

Keep Ethernet cables separated from power cables by at least 60mm where possible and never run them in the same channel of the rack without segregation.

- Category 6 (or better) Ethernet cable with RJ-45 jacks

1. Locate the cable to connect to the Netra X5-2 for Communications.
2. Plug the RJ-45 connector into the ILOM port.

Figure 2-5 Connecting to ILOM Port



3. Lead the cable neatly away from the rear panel.
4. Connect the other end of the cable to the LAN.
 - Refer to the Netra X5-2 for Communications hardware documentation for information on how to configure the Web browser application to connect to the console, and how to establish communications with the Netra X5-2 for Communications.

Software Installation - Netra and Server-based Platforms

Oracle Communications Session Delivery software requires software installation when deployed on Netra and Server-based platforms.

Installation Procedure

Software installation to Netra and server-based platforms includes the following high-level steps:

1. Ensure your device is set to boot from your boot media. This may be by way of USB or CD.
2. Insert your installation media in any USB slot or CD drive.
3. Power up the device, observing the boot cycle.
4. When power-up is complete, the device loads the Oracle Communications Session Delivery Product (OC-SDP) software. Wait for this to complete.
5. When notified, remove the boot media and allow the device to boot to the newly installed OC-SDP software.
(This step may not be required as some platforms support a boot priority mechanism that knows to boot from your hard drive after the installation is complete.)

Note:

Note that the OC-SDP boots by default to VGA (or as configured by BMC) during the installation. You can change this to serial temporarily during installation. After installation you can set the boot option to VGA or serial in the boot parameters. This setting is “permanent,” meaning that any device set to boot to VGA appears “dead” at serial (and vice-versa).

Logging Into the System

The Oracle Communications Session Delivery Product (OC-SDP) requires you to set passwords for the Admin and User accounts the first time you power up a new or factory reset system by way of local access. You cannot access the Admin and User accounts until you set the corresponding passwords. Use either an SSH or console connection when setting passwords. You log into your system after setting passwords.

Before you begin, plan your passwords to meet the following requirements:

- 8-64 characters
- Include three of the following:
 - Lower case letters
 - Uppercase letters
 - Numerals
 - Punctuation

The system leads you through the process for setting the Admin and User passwords, as follows:

1. Power up the SBC. The system prompts you to set the User account password.
2. At the prompt, type **acme**, and press ENTER. The system prompts you to enter the password that you want for the User account.
3. Type the User account password, and press ENTER.
4. Type **enable**, and press ENTER. The system prompts you to set the Admin account password.
5. Type **packet**, and press ENTER. The system prompts you to enter the password that you want for the Admin account.
6. Type the Admin account password, and press ENTER. The system logs you in as Admin.

First Steps after Software Installation

Oracle recommends the following steps after installation on the Oracle X7-2 platform.

1. Execute the OC-SDP **format hard-disk** command, per your requirements. See [Formatting the Disk Volume](#) for reference and instructions.
2. Turn off the OC-SDP using the **Halt** command. This provides you with a graceful software shutdown, after which the hardware is still powered on.
3. Power cycle the hardware using the power switch.

To configure the OC-SDP, refer to the *ACLI Configuration Guide*.

Boot parameter changes to consider prior to service configuration include:

- Set your "Target Name" to your preferred OC-SDP name.
- Verify your "Console Device", eg, com1 (serial).
- Set the "IP Address" to your preferred management port IP address.

 **Note:**

Note at the boot parameters that the default system is named “/boot/bzImage”. Be aware that upgrading code includes obtaining images with, for example, an SCz prefix and the .bz file extension.

Known Issues

Netra X5-2 for Communications

The **interface-mapping locate** command does not work with the Netra X5-2 for Communications onboard interfaces. The command does work with PCI interfaces installed on the platform.

The Onboard Ethernet ports of the Oracle X-series servers (X3-2, X5-2, and so on.) run natively at 10GBASE-T, which requires the use of Category 6a cabling. These ports negotiate down to 100BASE-T or 100BASE-T, but the negotiation might not succeed when you use incompatible cables. For example, do not use Cat5/5e cables or Cat6 cables not rated for 500MHz operation.

3

Oracle Server X7-2 Platform Preparation

Oracle Communications produces a variety of software products that run on the Oracle Server X7-2 platform, including Oracle session delivery applications.

Use your Hardware documentation to install and establish system management via ILOM. Then use the steps below to prepare the Oracle X7-2 for session delivery software installation.

Note:

The [ILOM Cable Connection procedure](#) also displays ILOM cabling.

1. Confirm applicable firmware on the server.
 - To check the firmware versions installed in the server, go to the Oracle Integrated Lights Out Manager (ILOM) web interface, and navigate to **System Information, Firmware**.
 - Software and firmware versions qualified for use with Oracle Session Delivery products include:
 - ILOM—v4.0.2.20.b
 - BIOS— 41.02.13.00
2. Upgrade or downgrade the firmware on the server as necessary. Go to https://docs.oracle.com/cd/E81115_01/index.html for ILOM upgrade instructions.
3. Configure the BIOS settings. (Settings navigation may differ based on the BIOS version.)
 - a. Observe the boot procedure, logged to the console during bootup, and use the documented key sequence to interrupt the boot and display the BIOS configuration dialogs. For example, pressing the F2 key is a common way to enter BIOS configuration from a terminal application that supports function keys.
 - b. Navigate to the Boot menu and, depending on the software distribution you are using, set the USB or CD as the first device followed by the disk controller. (Navigation: Boot)
 - c. Disable Hyper-Threading. (Navigation: Advanced, Processor Configuration, Hyper-Threading)
 - d. Disable CPU power limit. (Navigation: Advanced / CPU Power Management Configuration)
 - e. Disable C6 Reporting. (Navigation: Advanced / CPU Power Management Configuration, CPU C6 report)
 - f. Change Energy Performance to Performance. For example, set "ENERY_PERF_BIAS_CFG" mode to "PERF". (Navigation: Advanced / CPU Power Management Configuration, Energy Performance)


- g. To decrease boot up time, Oracle recommends disabling Intel PXE Boot Agent for both onboard and NIC ethernet ports. Press F2 and navigate to Advanced, Network Stack Configuration. Then disable IPv4 PXE support.

 **Note:**

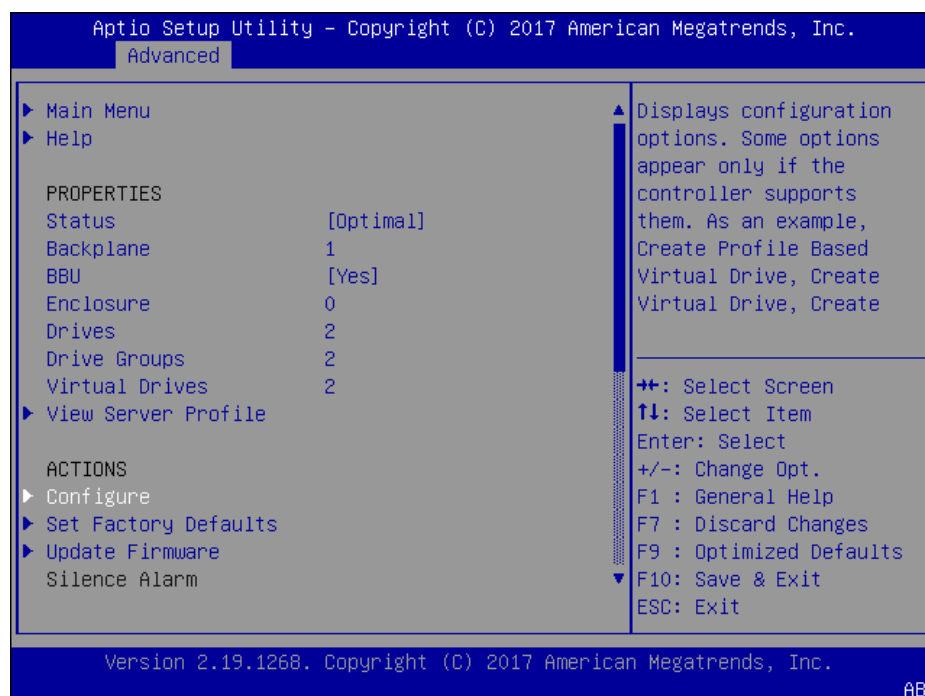
PXE boot is not supported in this release.

- h. Reboot the server.
4. Initialize the HDD.
- a. Open the ILOM remote system console to observe the system's boot cycle, and interrupt the boot cycle to enter the MegaRAID configuration utility.

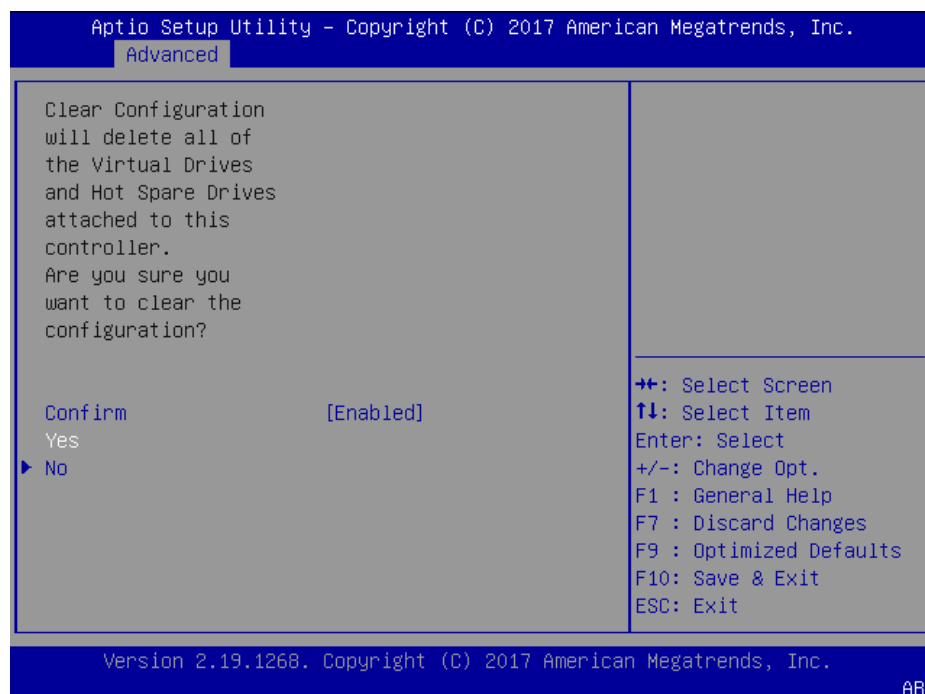
Figure 3-1 Selecting RAID Configuration




- b. Navigate the utility to establish your virtual drive's operation, initially including the **Configure** action.

Figure 3-2 Begin RAID Configuration


- c. Clear the configuration, regardless of the initial state.

Figure 3-3 Clear Any Existing RAID Configuration

- d. Access the menu from which you create a virtual drive.


Figure 3-4 RAID - Create Virtual Drive

- e. Set the RAID level to RAID-1.

Figure 3-5 Set Drive to RAID1


- f. Select your drives.

Figure 3-6 RAID - Select Drives




- g. It is common to select all drives at this point.


Figure 3-7 Select All Drives




- h. Save your RAID configuration.

Figure 3-8 Save RAID Configuration

- i. The system allows you to Confirm your configuration and continue with initialization.

Figure 3-9 Initialize RAID Configuration

- j. After initialization is complete, return to the Main Menu to Save and Exit.

Figure 3-10 Exit RAID Configuration

5. Perform a cold shutdown by removing all system power.

Cable the Oracle X7-2

After mounting the Oracle X7-2 in an equipment rack and installing all components, use the following instructions to connect all appropriate data cables to the ports before powering the system up and beginning the configuration.

Oracle has qualified the following configurations of the Oracle X7-2:

- Configuration A: One Four-port 10 GigE NIC
- Configuration B: Two Four-port 10 GigE NICs (each of the three slots are qualified)
- Configuration C: One QSFP NIC (in quad port mode only) and ONE Four-port 10 GigE NIC



Note:

The 40G interface speed is not supported.

On board interfaces for all configurations include:

- One RJ-45 serial management (SER MGT) port
- One 10/100/1000BASE-T RJ-45 Oracle Integrated Lights Out Manager (ILOM) service processor (SP) network management (NET MGT) port
- One 1000BASE-T RJ-45 Gigabit Ethernet (GbE) port, labeled NET 0
- Two 10/25GbE SFP+ Ethernet ports, labeled NET 1 and NET 2

- Two 10GBASE-T RJ-45 Gigabit Ethernet (GbE) ports, labeled NET 1 and NET 2



Note:

The 10/25GbE SFP+ Ethernet ports (NET 1 and NET 2) are the dedicated HA ports for the server. When using an SFP+ port, network connectivity is disabled on the 10GBASE-T RJ-45 GbE (NET 1 and NET 2) Ethernet ports.

Figure 3-11 Oracle X7-2 Configuration A (4x10 GigE NIC)




Figure 3-12 Oracle X7-2 Configuration B (Two 4x10 GigE NICs)





Figure 3-13 Oracle X7-2 Configuration B (One QSFP and One 4x10 GigE NICs)



▲ Caution:

Please review your Oracle X7-2 Product Notes. Notes for release 1.1.2 describes physical issues with some optical transceivers installed into an SFP28 port.

Oracle recommends using Category 6 (or better) for all Ethernet connections.

You do not need to use every port for proper operation.

Available Connections

Please read all of the information for each of the available connections prior to cabling the Oracle X7-2.

Port	Description	You Need:
NET (0-2)	<p>From left to right:</p> <ul style="list-style-type: none"> • 1 GigE ports - Net 0 • 10 GigE ports - Net 1, Net 2 <p>Enables you to connect the Netra X7-2 to your network.</p>	<p>A Category 6 (or better) Ethernet cable to connect to the NET 0 port to your network. Network parameters such as an IP address (can be provided by DHCP services or assigned a static address in the OS)</p> <p>Additional Category 6 (or better) Ethernet cables and Ethernet addresses as needed for additional connections to NET 0, 1 and 2.</p>
NET MGT	<p>Provides a 10/100/1000 BASE-T Ethernet connection to the Service Processor (SP) through an RJ-45 connector. The NET MGT port provides support connections to the SP using the Oracle Integrated Lights Out Manager (ILOM) CLI and Web interface. By default, the NET MGT port is configured to use DHCP to automatically obtain an IP address. Alternatively, you can assign a static IP address to the NET MGT port. To use the NET MGT port, you must configure its network settings. Once configured, use the NET MGT port IP address to log on to the device using a browser or secure shell.</p>	<p>Category 6 (or better) Ethernet cable to connect the NET MGT port to your network IP address for this port (required from DHCP or a static address)</p>

Port	Description	You Need:
SER MGT (COM1)	Provides a TIA/EIA-232 serial Oracle/Cisco standard connection to the SP through an RJ-45 connector. This interface connects to either Service Processor by default, but can be redirected to the host. Default settings: <ul style="list-style-type: none"> • 8N1: eight data bits, no parity, one stop bit • 9600 baud (change to 115200 baud) • Disable hardware flow control (CTS/RTS) • Disable software flow control (XON/XOFF) 	A terminal device (e.g., terminal, connection to a terminal server, or computer such as a laptop running terminal emulation software) A cable to connect the terminal device to the SER MGT (COM1) port
USB	Provides USB3.0 connection to the computer. The USB port is hot pluggable, so you can connect and disconnect USB cables without affecting server operations.	Installation media Note: Maximum USB cable length: 5 meters

Cable the Local Console

You can connect the Administration console to the local SER MGT (COM1) serial console port.

- To cable a serial console connection:
 - Serial console cable with an RJ-45 connector



Note:

Do not configure COM2 in the bootparams menu.

When configuring boot loader parameters, set the console to COM1 when you use SER MGT. The boot loader is accessible on all console ports, but only input from the active console port can be recognized by the Oracle X7-2.

1. Locate the appropriate cables to connect to the Oracle X7-2.
2. To cable a serial connection, insert the serial console cable into the SER MGT (COM1) port.

Figure 3-14 Connecting to USB and SER MGT (COM1) Ports



 **Note:**

Refer to the Oracle X7-2 hardware documentation for information on how to configure the terminal application to connect to the console, and how to establish communications with the Oracle X7-2.

3. For installation procedures, insert the USB stick in the USB port.
4. Lead the cables neatly away from the rear panel.
5. Plug in the cables to their respective destination components.

Connect ILOM to the Oracle X7-2

Use the following procedure to make a connection to the Oracle X7-2 Oracle Integrated Lights Out Manager (ILOM) port. For a remote permanent connection to the Service Processor over the ILOM connection, use the rear panel NET MGT port.


 **Note:**

Keep Ethernet cables separated from power cables by at least 60mm where possible and never run them in the same channel of the rack without segregation.

- Category 6 (or better) Ethernet

1. Locate the cable to connect to the Oracle X7-2 for Communications.
2. Plug the RJ-45 connector into the ILOM port.

Figure 3-15 Connecting to ILOM over the Network



3. Lead the cable neatly away from the rear panel.
4. Connect the other end of the cable to the LAN.
- Refer to the Oracle X7-2 hardware documentation (https://docs.oracle.com/cd/E72435_01/html/E72440/index.html) for information on how to configure the Web browser application to connect to the console, and how to establish communications with the Oracle X7-2.

Software Installation - Oracle X7-2 Platforms

The Oracle Communications Session Router (OCSR) requires software installation when deployed on the Oracle X7-2.

Install Procedure

Software installation to Oracle X7-2 includes the following high-level steps:

1. Insert your installation media in the USB slot. Alternatively, connect the ISO image via iLOM virtual media.

 **Note:**

Please review your Oracle X7-2 Product Notes. Notes for release 1.1.3 describes the requirement to maintain the default of SSL enabled for any OS installation.

2. Power on the Oracle X7-2 .
3. Observe the startup process and press F8 to enter the boot menu when it becomes available.
4. Select the bootable USB or ISO setting.

 **Note:**

You may need to scroll through down the list to reach the ISO setting.

5. Save and exit the boot menu. The Oracle X7-2 starts the OCSR installation.
6. Remove the USB media when prompted by the Oracle X7-2.
7. Allow the Oracle X7-2 complete its installation process and boot to the newly installed OCSR software.

The OCSR boots by default to VGA during the installation. The user can change this to serial "temporarily" during install. Within the boot parameters and after installation, however, the user can set the boot option to VGA or serial. This setting is "permanent", meaning that any device set to boot to VGA appears "dead" at serial (and vice-versa).

 **Note:**

There is no physical VGA on the Oracle X7-2, but VGA emulation is available via the iLOM remote console.

Logging Into the System

The Oracle Communications Session Delivery Product (OC-SDP) requires you to set passwords for the Admin and User accounts the first time you power up a new or factory reset system by way of local access. You cannot access the Admin and User accounts until you set the corresponding passwords. Use either an SSH or console connection when setting passwords. You log into your system after setting passwords.

Before you begin, plan your passwords to meet the following requirements:

- 8-64 characters
- Include three of the following:
 - Lower case letters
 - Uppercase letters
 - Numerals
 - Punctuation

The system leads you through the process for setting the Admin and User passwords, as follows:

1. Power up the OCSR. The system prompts you to set the User account password.
2. At the prompt, type **acme**, and press ENTER. The system prompts you to enter the password that you want for the User account.
3. Set your User account password, and press ENTER.
4. Type **enable**, and press ENTER. The system prompts you to set the Admin account password.
5. Type **packet**, and press ENTER. The system prompts you to enter the password that you want for the Admin account.
6. Set your admin account password, and press ENTER. The system logs you in as Admin.

 **Note:**

Setting passwords is also covered in the *ACLI Configuration Guide*.

First Steps after Software Installation

Oracle recommends the following steps after installation on the Oracle X7-2 platform.

1. Execute the Oracle Communications Session Delivery Product's **format hard-disk** command, per your requirements. See [Formatting the Disk Volume](#) for reference and instructions.
2. Turn off the Oracle Communications Session Delivery Product using the **Halt** command. This provides you with a graceful software shutdown, after which the hardware is still powered on.
3. Power cycle the hardware using the power switch, a power controller or by physically disconnecting and reconnecting the power cable.

To configure the Oracle Communications Session Delivery Product, refer to the *ACLI Configuration Guide*.

Boot parameter changes to consider prior to service configuration include:

- Set your "Target Name" to your preferred Oracle Communications Session Delivery Product name.
- Verify your "Console Device", eg, com1 (serial).
- Set the "IP Address" to your preferred management port IP address.
- Set the "Netmask" for your management port IP address.
- Set the "Gateway" address for your management port IP address.

 **Note:**

Note at the boot parameters that the default Boot File is "/boot/bzImage". Be aware that upgrading code includes obtaining images with, for example, an SCz prefix and the .bz file extension.

Known Issues

Oracle X7-2 for Communications

The **interface-mapping locate** command does not work with the Oracle X7-2 for Communications onboard interfaces. The command does work with PCI interfaces installed on the platform.

The Onboard Ethernet ports of the Oracle X-series servers (X3-2, X5-2, X7-2 and so on.) run natively at 1GBASE-T or 10GBASE-T, which requires the use of Category 6a cabling. These ports negotiate down to 1000BASE-T or 100BASE-T, but the negotiation might not succeed when you use incompatible cables.

4

Session Delivery Products as Virtual Machines

You can deploy version S-CZ8.1.0 of Oracle's Session Delivery Products as Virtual Machine (VM). This document refers to Session Delivery Products generically. See your software version's Release Notes to verify your product's support for deployment as a virtual machine.

VM deployment types include:

- A standalone (not orchestrated) instance Oracle Communications Session Delivery Product operating as a virtual machine running on a hypervisor, and
- Virtual Machine(s) deployed within an Orchestrated Network Functions Virtualization (NFV) environment.

Standalone OC-SDP VM deployment instances are always supported. Support within an orchestrated environment is dependent on orchestrator and OC-SDP version. High Availability configurations are supported by both deployment types.

Hypervisor platform and resource requirements are version-specific. Refer to your *Release Notes* for applicable requirements, recommendations, caveats for qualified platforms.

Hyper-threading and CPU Affinity

Due to the polling operation of DPDK, using Hyper-threaded cores can degrade the OC-SDP's performance. Oracle recommends the user disable Hyper-threading in the host system if possible, or configure CPU affinities on the hypervisor to ensure mapping from only one virtual CPU to each physical CPU core.

The user can learn how to configure CPU affinity via their hypervisor documentation.

Host Hypervisor CPU Pinning

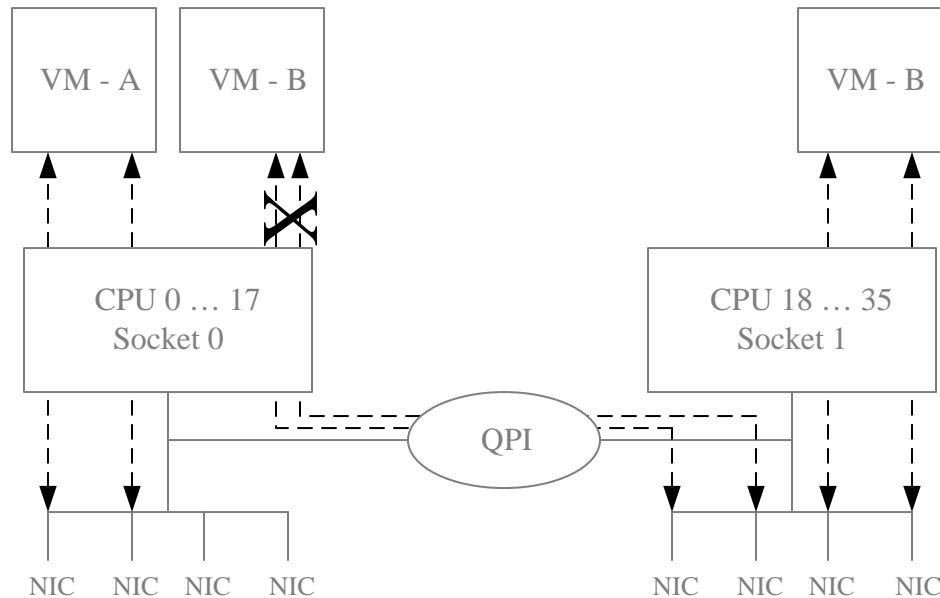
Many hardware platforms have built in optimizations related to VM placement. For example, some CPU sockets may have faster local access to Peripheral Component Interconnect (PCI) resources than other CPU sockets. Users should ensure that VMs requiring high media throughput are optimally placed by the hypervisor, so that traversal of cross-domain bridges, such as QuickPath Interconnect (QPI), is avoided or minimized.

Some hypervisors implement Non-Uniform Memory Access (NUMA) topology rules to automatically enforce such placements. All hypervisors provide manual methods to perform CPU pinning, achieving the same result.

The diagram below displays two paths between the system's NICs and VM-B. Without configuring pinning, VM-B runs on Socket 0, and has to traverse the QPI to access

Socket 1's NICs. The preferred path pins VM-B to Socket 1, for direct access to the local NICs, avoiding the QPI.

Figure 4-1 Contrast of Data Paths with and without CPU Pinning



The user can learn how to pin CPUs via their hypervisor documentation.

Configuration Overview

Oracle Communications Session Delivery Product Virtual Machine (VM) deployments require configuration of the VM environment and, separately, configuration of the OC-SDP itself. VM-specific configuration on the OC-SDP includes boot parameter configuration, enabling functionality and performance tuning.

During VM installation, you can configure operational information within the system's boot parameters, including:

- IP address
- Host name

Refer to the instructions on installing the OC-SDP on each supported hypervisor to see how and when you do this during installation. Refer to the *ACLI Configuration Guide* for instruction on changing these settings after installation.

During VM installation, the OC-SDP sets default functionality, assigning cores to signaling and media forwarding. If you need DoS and/or transcoding functionality, configure the applicable cores after installation and before performing service. See the *ACLI Configuration Guide* for core configuration instructions.

Applicable performance tuning configuration after deployment includes:

- Media manager traffic/bandwidth utilization tuning
- Datapath-related CPU core allocation

See the *ACLI Configuration Guide* for instructions on VM tuning configuration.

 **Note:**

For Xen-based hypervisors, the default boot mode uses DHCP to obtain an IP address for the first management interface (wancom0) unless a static IP is provisioned. Note that DHCP on wancom0 does not support lease expiry, so the hypervisor must provide persistent IP address mapping. If persistent IP address mapping is not provided, the user must manually restart the VM whenever the wancom0 IP address changes due to a manual change or DHCP lease expiry.

Beyond installation, VM-related functional support, and VM-related tuning, you perform basic OC-SDP configuration procedures after installation, including:

- Setting Passwords
- Setup Product
- Setup Entitlements
- Service configuration

Refer to the *ACLI Configuration Guide* for instructions on these procedures.

Refer to the documentation listed in the *About This Guide* section of this document for other OC-SDP-related configuration.

VLAN Support

Version S-CZ8.1.0 supports VLANs within all supported virtual machine environments and interface architectures. Refer to the *ACLI Configuration Guide* for instructions on configuring VLANs on the OC-SDP. Note that when you configure a VLAN, the OC-SDP requires VLAN tags to be included in the packets delivered to and from the VM.

Oracle recommends that you evaluate the VLAN support of your deployment's hypervisor and interface I/O mode before implementation to ensure secure support for the transmission and receiving of VLAN-tagged traffic. Please consult your hypervisor's vendor documentation for details.

Provisioning Entitlements

VNF products licensing follows the standard C-series self-entitlements licensing model. Refer to the *ACLI Configuration Guide* for instructions on setting entitlements.

5

Virtual Machine Platforms

Oracle distributes virtual machine templates, each containing a virtual disk image and default configuration for the supported profile of each VM platform. VM platform support is dependent on your Oracle product version.

This section addresses requirements associated with running applicable software as virtual machines. It also provides basic instructions on loading and starting machine templates.

VM distributors maintain extensive documentation sites. You must use those vendors' documentation for full explanations and instructions on VM deployment and operation.

Create and Deploy on Oracle VM Manager

This section provides detail on deploying Oracle Communications Session Delivery products in an Oracle Virtual Machine environment and assumes Oracle VM Manager 3.4.2. The section also assumes the user understands deployment in these environments and that the majority of deployment tasks, from hardware installation and startup to VM resource and management setup, is complete.

For information on Oracle OVM, including server and manager documentation, refer to the following links. The bottom link opens the page specifically to Oracle OVM version 3.4.2.

<http://www.oracle.com/technetwork/documentation/vm-096300.html>
http://docs.oracle.com/cd/E64076_01/

Once hardware, VM resources and VM management are in place, the user prepares the VM management environment. High level steps include:

- Discover Oracle VM Servers
- Discover Storage (File Server or SAN)
- Create a Virtual Machine Network
- Create a Server Pool
- Create a Storage Repository

Note:

The following procedure describes a typical deployment. The system may display different screens, depending on the deployment.

Oracle Communications Session Delivery product-specific setup steps include the following.

- Add Resources to Storage Repository

- Import an assembly
- Import a virtual machine template (from the assembly)
- Create a virtual machine from a template
- Configure processing resources
- Assign Networks
- Specify Boot Order
- Connect to the console
- Start your virtual machine

Use the Oracle VM Manager to deploy your VMs. Browsing the manager displays the management application, with tabs across the top for VM Manager configuration and display, and two panes providing access to controls and components for specific VM management.

Follow the steps below to deploy your VM(s):

1. From the Oracle VM Manager application Home page, navigate to the **Repositories** tab.
2. Click the **Virtual Appliances** folder on the left pane, then click the download icon from the center pane.

Oracle VM Manager displays the **Import Virtual Appliance** dialog.

3. Type the URL, either the HTTP or the FTP server, of your .ova appliance into the **Virtual Appliance download location** text box.
4. Check the **Create VM** checkbox.
5. Select the server pool to which your new machine belongs from the **Server Pool** dropdown.
6. Click OK.

Oracle VM manager creates your Virtual Machine.

7. Select **VM Files** folder on the left pane and verify the file is present.
8. Select the **Servers and VMs** tab and select the server pool.
9. Click on the **Perspective** drop down menu and select **Virtual Machines**.
10. Right click on the Virtual Machine and select **Edit**.
11. Edit your VM to configure the desired CPU, memory and core resources. Consider the following settings:
 - Operating System - Oracle Linux 7
 - Domain Type - XEN HVM PV Drivers
 - Memory - Set according to your deployment (defaults to 8G)
 - Processors - Set according to your deployment (defaults to 4)
12. Open the **Networks** tab to manage your networks using the **Network** drop-down selection box. OVM does not display MAC addresses until the user applies the configuration.

The OC-SDP enumerates and binds network interfaces in the order presented by the hypervisor to the virtual machine. If it presents 3 or less interfaces, the bind

order is wancom0, s0p0, s1p0. If it presents more than 3 interfaces, the bind order is:

- a. wancom0
- b. wancom1
- c. wancom2
- d. spare
- e. s0p0
- f. s1p0
- g. s0p1
- h. s1p1

If your hypervisor randomly allocates addresses for network interfaces, the interface order at the hypervisor does not necessarily match that at the OC-SDP. You can use the **interface-mapping show** command to determine the MAC address to interface order, and if necessary, adjust it using the **interface-mapping swap** command.

13. If you want to increase the default disk size, click the **Disks** tab and the pencil icon to set a larger disk size.

Do not decrease disk size.

14. Click the **Boot Order** tab and ensure that Disk is the first (or only) option in the boot order list.

15. Click OK.

The system applies your settings to your VM.

16. Click the **Console** icon from the menu bar.

Oracle VM Manager displays a terminal screen with the serial CLI operational.

17. Highlight the target and click the **Start** button.

Oracle VM Manager starts your VM, displaying the startup sequence and, ultimately, providing ACLI access in the console.

Create and Deploy on KVM

For complete KVM documentation, refer to <http://www.linux-kvm.org/page/Documents>.

1. Install the Virtualization Host group and virt-install.

```
# yum groupinstall "Virtualization Host"  
# yum install virt-install
```

2. Extract the image.

```
# tar xvf nnSCZ739.64-img-vm_kvm.tar  
nnSCZ739.64-img-vm_kvm.ovf  
nnSCZ739.64-img-vm_kvm.qcow2  
legal.txt
```

3. Use virt-manager to create the management and media network interfaces.

- Create a virtual bridged network for management interfaces.
 - Create virtual networks for media interfaces.
4. Provision a new virtual machine.

```
# virt-install \
  --name SBC739 \
  --description "nnSCZ739 KVM" \
  --os-type=Linux \
  --os-variant=rhel7 \
  --ram=8192 \
  --vcpus=4 \
  --disk path=/opt/nnSCZ739.64-img-
vm_kvm.qcow2,bus=virtio,size=10,format=qcow2 \
  --network bridge=br-Mgmt \
  --network bridge=br-Mgmt \
  --network bridge=br-Mgmt \
  --network bridge=br-Mgmt \
  --network network=medial \
  --network network=media2 \
  --import \
  --cpu host
```

 **Note:**

Use interface-mapping to pin the four br-Mgmt network interfaces to wancom0, wancom1, wancom2, and spare.

--name

Identify a unique name for the virtual machine on this hypervisor.

--description

Describe this virtual machine.

--os-type

Specify the operating system type.

--os-variant

Optimize the configuration for a specific operating system.

--ram

Allocate a specific amount of RAM to this virtual machine.

--vcpus

Allocate a specific number of virtual CPUs to this virtual machine.

--disk

Specify the path to the disk image.

--network

Connect the virtual machine to a host network.

--import

Skip the operating system installation process and build a guest around the disk image specified with `--disk`.

--cpu

Configure the CPU model and CPU features exposed to the virtual machine.

See `man virt-install` for more information.

 **Note:**

The `--cpuset` and `--numatune` options may be used to establish CPU affinity and socket pinning.

Create and Deploy on VMware®

This section provides detail on deploying Oracle Communications Session Delivery products over the ESXi hypervisor and assumes VMware 6. The section also assumes the user understands deployment in these environments and that the majority of deployment tasks, from hardware installation and startup to VM resource and management setup, is complete.

For information on VMware 6, which is also supported, refer to the following link.

<https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-6-pubs.html>

Before You Begin:

- Confirm that the VMware version 6 Hypervisor is installed on an appropriate network server.
- Confirm that the server has 40GB of space for this installation.

 **Note:**

The following procedure describes a typical deployment. The system may display different screens, depending on the deployment.

Detail on Oracle Communications Session Delivery product-specific setup steps is shown below.

1. On the vSphere Client application Home page, go to **File, Deploy OVF Template File**.
2. On the Source screen, browse to the target .ova file, and click Next.
3. On the End User License Agreement screen, click Accept and click Next.
4. On the Name and Location screen, do the following and click Next.
 - Name. Enter a name for the template.
 - Inventory Location. Select the location where you want to deploy the template.
5. On the Host / Cluster screen, select the host or cluster where you want to run the deployed template, and click Next.

6. If applicable to your deployment, select the resource, from the Resource Pool screen, in which you want to deploy this template, and click Next.
7. On the Storage screen, select the destination storage for the virtual machine files, and click Next.
8. On the Disk Format screen, select Thick Provisioned Lazy Zeroed, and click Next.
9. On the Network Mapping screen, map the networks used in this template to networks in your inventory, and click Next.


The OC-SDP enumerates and binds network interfaces in the order presented by the hypervisor to the virtual machine. If 3 or less interfaces are presented, the bind order is wancom0, s0p0, s1p0. If more than 3 interfaces are presented, the bind order is:

- a. wancom0
- b. wancom1
- c. wancom2
- d. spare
- e. s0p0
- f. s1p0
- g. s0p1
- h. s1p1

If your hypervisor randomly allocates addresses for network interfaces, the interface order at the hypervisor does not necessarily match that at the OC-SDP. You can use the **interface-mapping show** command to determine the MAC address to interface order, and if necessary, adjust it using the **interface-mapping swap** command.

10. On the properties screen, enter parameters per your configuration, and click Next.

Figure 5-1 Properties Dialog for OVF Deployment



 **Note:**

Do not enter plaintext values in the password fields.

11. On the Ready to Complete screen, review the selections, and do one of the following:
 - Edit a previous screen by way of the Back button.
 - When the edits are completed, or if no edits are required, click Finish
 The system displays the Deploying screen.

Figure 5-2 OVF Deployment Status Dialog



When the Deploying screen closes, VM creation and deployment is completed.

Create and Deploy on Hyper-V®

Follow the steps below to deploy the Oracle Communications Session Delivery Product (OC-SDP) on Hyper-V on Windows 2012 R2 (Generation 1). This procedure assumes you understand deployment with Hyper-V hypervisor and that the majority of deployment tasks, from hardware installation and startup to Virtual Machine (VM) resource and management setup, are complete.

For information on Hyper-V, refer to the following link.


<https://docs.microsoft.com/en-us/windows-server/virtualization/virtualization>

Before You Begin:

- Refer to your OC-SDP version's Release Notes for minimum required memory and CPUs.
- Confirm that the Hyper-V hypervisor is installed on an appropriate network server.
- Confirm that the server has 40GB of space for this installation.
- Confirm the number of network interfaces needed for your deployment. (Wancom0, wancom1 and wancom2 should be Legacy Network Adapters; all others should be Network Adapters (PV).
- Confirm the amount of memory needed for your deployment.
- Confirm the number of processors to use for your deployment.
- Confirm your .vhd (Virtual Hard Drive) file is available to Hyper-V in a permanent location. Keeping this Oracle distribution on the same physical server as the Hyper-V manager ensures the best access to it during OC-SDP operation.

The following procedure describes an example that provides basic deployment guidelines. Steps for deploying your system may differ. You may, for example, decide not to use the wizard. In addition, the Hyper-V Manager provides access to its controls and wizards from multiple entry points.

Instances of Hyper-V Manager may display the **Actions** dialog in the upper-right pane of the manager's main window, but you may find your **Actions** controls elsewhere in your manager. Regardless of access, you use this dialog to start and run the **New, Virtual Machine** wizards used in this procedure.



1. Start the Hyper-V Manager.
2. Start the **Virtual Switch Manager** from the **Actions** dialog.
Hyper-V displays the **Virtual Switch Manager** dialog.
3. Click the **Create Virtual Switch** button.
Hyper-V modifies the **Virtual Switch Manager** dialog, presenting fields within which you specify your new switch.
4. Add virtual networks for each management and media interface. Set the following on the **Create Virtual Switch** dialog for each switch you create:
 - Virtual switch **Name**
 - Uncheck **Allow management operating system to share this network adapter**
 - The switch for wancom1, wancom2 can be **internal** or **external**. High availability via external (eg, separate hypervisor platforms) is preferred.
 - All other switches must be **external**.

During installation, the OC-SDP enumerates and binds network interfaces in the order presented by the hypervisor to the OC-SDP. This "presented" order is the same order in which you create networks. If the manager presents 3 or less interfaces, the bind order is wancom0, s0p0, s1p0. If it presents more than 3 interfaces, the bind order is:

- a. wancom0
- b. wancom1
- c. wancom2
- d. spare
- e. s0p0

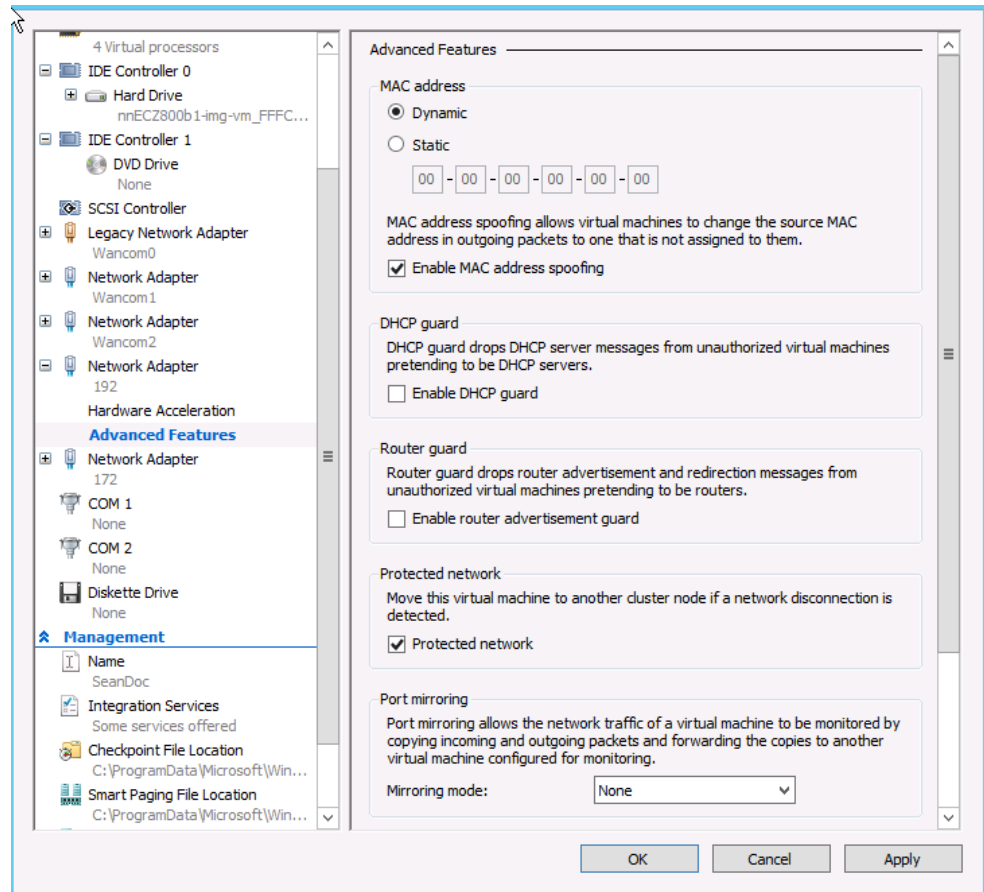
- f. s1p0
 - g. s0p1
 - h. s1p1
5. Click **New, Virtual Machine**
Hyper-V displays the introductory page of the **New Virtual Machine Wizard**.
6. Click **Next**.
Hyper-V advances through the **New Virtual Machine Wizard** pages each time you click **Next**. The wizard allows you to go back to the Previous page, Cancel the wizard, or Finish the wizard with the respective buttons. Your procedure through the wizard may vary, depending on your infrastructure and intent.
7. Enter or select at least the following as you progress through the wizard.
 - a. Type a name for your VM in the **Name** field.
 - b. Select **Generation 1** as your machine type.
 - c. Assign the desired memory.
 - d. Click **Next**, skipping the Configure Networking dialog. You add networks later in the process.
 - e. Connect to the Virtual Hard Disk you downloaded by selecting the **Use an existing virtual hard disk** radio button and browsing to your .vhd file.
 - f. Finish

The Hyper-V Manager returns to the main dialog, displaying your new machine in the Virtual Machine list.
8. Right click the VM.
The Hyper-V Manager displays a pop-up menu.
9. Click **Settings**
The Hyper-V Manager displays the **Settings** dialog for your Virtual Machine displaying the **Add Hardware** controls in the right-hand pane.
10. Select **Legacy Network Adapter** and click the **Add** button.
The Hyper-V Manager displays the **Legacy Network Adapter** dialog.
11. Select wancom0 you configured for your Virtual Machine from the drop-down listbox and click the **Apply**, then the **OK** buttons. Repeat this step for wancom1 and wancom2 if you are using these interfaces.

The Hyper-V Manager returns to the initial Settings dialog and adds this adapter to your machine's component list in the left-side pane. Only configure wancom0, wancom1 and wancom2 as **Legacy Network Adapters**.
12. Select **Network Adapter** and click the **Add** button.
The Hyper-V Manager displays the **Network Adapter** dialog.
13. Select the first adapter after wancom0 that you configured for your Virtual Machine from the drop-down listbox and click the **Apply**, then the **OK** buttons.
14. Repeat the previous step for the rest of your adapters, referring to the order described above.
15. Re-Open the **Settings ...** dialog.

16. For each **Network Adapter**, click the + sign beneath it to display the **Hardware Acceleration Advanced** links.
17. For each **Network Adapter's**, **Hardware Acceleration** settings, uncheck the **Enable virtual machine queue** checkbox. **Apply** and **OK** these changes.
18. For each media interface's **Network Adapter, Advanced** settings, check the **Enable MAC address spoofing** checkbox. **Apply** and **OK** these changes.

Figure 5-3 VM Settings on Hyper-V



19. Select **Processor** from the left-side pane, increase the number of processors for your deployment and click **Apply**, then **OK** to close the VM **Settings** dialog.
20. Right click your OC-SDP VM and Click **connect**.
The Hyper-V Manager displays a VM connection dialog.
21. Click the **Power Button** icon to turn on your OC-SDP VM.
22. Observe the machine boot process via the connection window until the boot finishes.

 **Note:**

If your hypervisor randomly allocates addresses for network interfaces, the interface order at the hypervisor may not match that at the OC-SDP. If necessary, you can use the **interface-mapping show** command to determine the MAC address-to-interface order and adjust it using the **interface-mapping swap** command.

23. Proceed with OC-SDP configuration.

Configure OpenStack Heat Template

The Oracle Communications Session Delivery Product supports Heat templates when launching virtual machines in OpenStack. Heat is OpenStack's orchestration service, and a Heat Orchestration Template (HOT) is a YAML file that defines the networks, security group, and other resources available for each virtual machine. During orchestration, Heat can simultaneously launch multiple virtual machines that work together as HA pairs.

Before editing the Heat template for the Oracle Communications Session Delivery Product:

- Check that your OpenStack version is Newton or newer.
- Confirm your OpenStack administrator has installed the Orchestration service Heat.
- Install the OpenStack CLI client on a machine that can access your OpenStack services.
- Set environmental variables using the OpenStack RC file.
- Verify an OC-SDP image has been uploaded to OpenStack.
- Gather the following information from your OpenStack administrator:

Information	GUI Path	CLI command
Networks	From the Project tab, click Network and then Networks .	<code>openstack network list</code>
Images	From the Project tab, click Compute and then Images .	<code>openstack image list</code>
Flavors	From the Project tab, click Compute and then Instances . From the Launch Instance window, click Flavor . Click Cancel to close the Launch Instance window.	<code>openstack flavor list</code>
Security Groups	From the Project tab, click Network and then Security Groups .	<code>openstack security group list</code>

Refer to the Heat Orchestration Template specification at openstack.org.

1. Extract the compressed Heat Orchestration Template (HOT) file bundled with your software package.

```
tar xzf <filename>_HOT.tar.gz
tar xf <filename>_HOT_newton.tar
```

This creates a local directory called `newton`.

2. Find the environment file.

The HA environment file is `newton/properties/sdHaParams.yaml`.
The standalone environment file is `newton/properties/sdStandaloneParams.yaml`.

3. Configure the following required parameters:

 **Note:**

Some parameters are only found in the HA environments.

Parameter	Description	HA only
<code>primaryName</code>	Displays in the OpenStack GUI in the Instance Name column. The value gets passed in the bootparams as the Target Name.	No
<code>secondaryName</code>	Displays in the OpenStack GUI in the Instance Name column. The value gets passed in the bootparams as the Target Name.	Yes
<code>highAvailability</code>	Enable or disable HA for this template. Always set to <code>true</code> in the HA environment file and <code>false</code> in the standalone environment file.	No
<code>flavor</code>	The name or ID of the OpenStack flavor.	No
<code>image</code>	The name or ID of the previously uploaded OC-SDP image.	No
<code>availabilityZone</code>	Specify the availability zone where the OC-SDP will be placed.	No
<code>securityGroup</code>	Specify the security group. The default security group drops incoming traffic, so either change the parameters of the default security group or create a new security group that allows incoming traffic.	No
<code>affinityPolicy</code>	Set the affinity policy for an HA pair. Affinity policy controls whether the nodes of an HA pair run on the same or separate physical hosts. Possible values are: <ul style="list-style-type: none"> • <code>anti-affinity</code> (default)—The nodes run on separate physical hosts. • <code>affinity</code>—The nodes run on the same physical host. • <code>none</code>—no affinity policy is applied 	Yes
<code>userPass</code>	The SHA1 hash of the user passphrase. Use the following Linux command to generate the SHA1 hash of the passphrase:	No

```
echo -n "<passphrase>" | shasum
```

where `<passphrase>` is the passphrase you want to use.

Parameter	Description	HA only
adminPass	The SHA1 hash of the admin password. Use the following Linux command to generate the SHA1 hash of the passphrase: <pre>echo -n "<passphrase>" sha1sum</pre> where <passphrase> is the passphrase you want to use.	No
networks	Specify the name of the networks to attach to each network interface. Media networks must be previously created, but private tenant networks for wancom1 and wancom2 can be automatically created with the "auto" value. Blank entries mean the network interface is not attached to a network.	No
portSecurityPolicy	Specify the port security policy currently in use for each network. Possible values are: <ul style="list-style-type: none"> net-on—Ports on this network have Port Security enabled by default net-off—Ports on this network have Port Security disabled by default port-off—Heat requests port security be disabled when creating ports on this network The network's port security policy defines the NeutronPort's policy, but privileged users can override the policy for individual ports with the <code>port-off</code> option.	No
primary:fixed_ips	Specify the IP addresses for the network interfaces on the primary HA pair. The wancom0 address gets passed in the bootparams as IP Address.	No
secondary:fixed_ip	Specify the IP addresses for the network interfaces on the secondary HA pair. The wancom0 address gets passed in the bootparams as IP Address.	Yes
virtualIPs	Specify the virtual IP addresses for HA media interfaces.	Yes
virtualMACs	Specify the virtual MAC addresses for HA media interfaces.	Yes

 **WARNING:**

Stack creation will fail under the following circumstances:

- The value for portSecurityPolicy is incorrect.
- The IP address is not available on the network.
- The user does not have permission to specify IP addresses on the network.

4. Use the CLI to create the stack.

In the arguments, specify:

- the template file
- the environment file
- the name of the stack

For example:

```
openstack stack create -t newton/stacks/SD.yaml -e newton/  
properties/sdHaProperties.yaml haSBC
```

6

Boot Management

Boot Management includes the tasks needed to ensure the system is operating according to the users requirements as it starts up. Requirements met by properly managing system boot include defining management access IP, specifying the load to boot and specifying a system name. The user may set this information manually or configure the operational environment to provide it.

Boot management consists of tasks working with the following:

- **Boot Loaders**—The user needs to perform file management tasks to ensure that the software used to boot the system is compatible with the application system software itself. This typically includes verifying boot loader and application system software version for compatibility and placing the correct boot loader software in the correct location.
- **Boot Parameters**—The user sets boot parameters to specify their requirements for boot, including defining management access IP, specifying the load to boot and specifying a system name.
- **Boot Flags**—The user can, optionally, set special boot parameters called boot flags to further define how the system boots. The user may also set boot flags for diagnostic purposes under the guidance of Oracle support personnel.

Boot Loader Overview

Boot loader software loads the application to run on a platform. As such, boot loader software must be correct before system startup. Oracle Communications Session Delivery product distributions include and install the correct boot loader during application installation, meaning you need not consider boot loader version during first installation procedures. Application software upgrades do not update boot loaders. For this reason, you need to verify this compatibility manually. The following topics provide information about identifying the need and performing the updates.

Stage3 Boot Loader

Every new software release includes a system software image and a Stage3 boot loader. Oracle recommends you update this boot loader with every software upgrade, as described in the Software Upgrade section. Be sure to perform this update before booting the new system image.

The Stage3 boot loader is generally backward compatible with previous releases, but Oracle recommends that the Stage3 boot loader be installed from the same Major.Minor version as the system image. It is not normally necessary to update the boot loader when installing a maintenance or patch release when the Major.Minor release is the same.

For example, the same nnSCZ720.boot can be used with S-CZ720, S-CZ720m1, and so forth system software. But it should be upgraded when installing S-CZ730 system software to match that Major.Minor release.

The boot loader file name corresponds to the software image filename. For example, if the software image filename is nnECZ720.64.bz, the corresponding Stage3 boot loader filename is nnECZ720.boot. The Stage3 boot loader is compatible with previous releases.

Stage 3 boot loader upgrade procedure can be found in the [Update the Stage 3 Bootloader](#) section of this guide.

**Note:**

The OC-SDP does not support uploading the boot loader by way of the Web GUI.

Boot Parameters

Boot parameters specify the information that your device uses at boot time when it prepares to run applications.

This section explains how to view, edit, and implement device's boot parameters, and boot flags. Boot parameters:

- Allow you to set the IP address for the management interface (wancom0).
- Allow you to set a system prompt. The target name parameter also specifies the title name displayed in your web browser and SNMP device name parameters.
- Specify the software image to boot and from where the system boots that image.

**Note:**

You must configure all three components of an IPv6 address, including address, mask and gateway, in your system's boot parameters for wancom0 addressing. Configure the mask as a forslash (/) after the address followed by the mask in number of bits. The system requires all three components for IPv6 Neighbor Discovery to work properly.

Boot flags are arguments to a specific boot parameter, and allow functional settings, such as the use of DHCP for acquiring a management port address, as well as various diagnostic startup configurations.

Configuring boot parameters has repercussions on your system's physical and network interface configurations. When you configure these interfaces, you can set values that might override the boot parameters.

The bootparam configuration list is shown below.

```
[Acme Boot]: p
Boot File      : /boot/bzImage
IP Address     : 172.44.12.89
VLAN          :
Netmask       : 255.255.0.0
Gateway       : 172.44.0.1
IPv6 Address   : 3fff:ac4:6001:0:208:25ff:fe05:f470/64
```

```

IPv6 Gateway      : 3fff:ac4:6001::ac4:6001
Host IP           :
FTP username      :
FTP password      :
Flags             : 0x00000040
Target Name       : ORACLE
Console Device    : COM1
Console Baudrate  : 115200
Other             :

```

```

[Acme Boot]: ?
?              - print this list
@              - boot (load and go)
p              - print boot params
c              - change boot params
v              - print boot logo with version
r              - reboot
s              - show license information

```

Boot Parameter Definitions

The system displays all boot parameters when the user configures them after a boot interrupt. The system hides some boot parameters from the ACLI because the user should not configure them. If changed improperly, these parameters can cause the system to fail.

The following table defines each of the parameters that are visible when the user configures after a boot interrupt.

Boot Parameter	Description
Boot File	The name and path of the software image you are booting. Include the absolute path for a local boot from the local /boot volume and for a net boot when a path on the FTP server is needed.
IP Address	IP address of wancom0.
VLAN	VLAN of management network over which this address is accessed. Note: VLANs over management interfaces are not supported on the Acme Packet 4600 And Acme Packet 6xxx.
Netmask	Netmask portion of the wancom0 IP Address.
Gateway	Network gateway that this wancom0 interface uses.
IPv6 address	Version 6 IP address/mask of wancom0. Configure the mask as a forslash (/) after the address followed by the mask in number of bits.
IPv6 Gateway	Version 6 network gateway that this wancom0 interface uses.
Host IP	IP Address of FTP server from which to download and execute a software image.
FTP Username	FTP server username
FTP password	FTP server password

Boot Parameter	Description
Flags	Codes that signal the system from where to boot. Also signals the system about which file to use in the booting process. This sequence always starts with 0x (these flags are hexadecimal).
Target Name	Name of the Oracle Communications Session Delivery Product as it appears in the system prompt. For example, ORACLE> or ORACLE#. You need to know the target name if you are setting up an HA node. This name is required to be unique among Oracle Communications Session Delivery Products in your network. This name can be 63 characters or less.
Console Device	Serial output device type, dependent on platform. COM1 applies to virtual serial consoles, VGA to virtual video console. VGA is the default on VMware and KVM. COM1 is the default on OVM.
Console Baud Rate	The speed in bits per second which the console port operates at. It operates at 115200 BPS, 8 data bits, no stop bit, parity NONE.
Other	Allows miscellaneous and deployment-specific boot settings.

Boot Flags

Boot flags enable system boot behavior(s). The user can set a single flag, or add hex digits to set multiple flags.

- 0x00000008 Bootloader ~7 seconds countdown
- 0x00000040 Autoconfigure wancom0 via DHCP enable - VM platforms only
- 0x00000080 Use TFTP protocol (instead of FTP) enable - VM platforms only
- 0x00000100 Bootloader ~1 seconds quick countdown - VM platforms only

The following boot flags should only be used as directed by Oracle support:

- 0x00000001 acme.ko network module security override
- 0x00000002 Kernel debug enable
- 0x00000004 Crashdump disable
- 0x00000010 Debug sshd enable
- 0x00000020 Debug console enable getty
- 0x00001000 Userspace debug enable
- 0x00100000 Uniprocessor enable (SMP disable)
- 0x20000000 Fail-safe boot enable
- 0x40000000 Process startup disable (flatspin mode)

Never enter any other values without the direction of Oracle support. Some diagnostic flags are not intended for normal system operation.

Changing Boot Parameters

You can access and edit boot parameters by using either the ACLI or by interrupting the system boot process.

 **Note:**

Changes to boot parameters do not go into effect until you reboot the system.

Change Boot Parameters from the ACLI

To access and change boot parameters from the ACLI:

1. In Superuser mode, type `configure terminal`, and press Enter.

```
ORACLE# configure terminal
```

2. Type `bootparam`, and press Enter. The boot device parameters display.

```
ORACLE(configure)# bootparam  
'.' = clear field; '-' = go to previous field; q = quit  
Boot File      : /boot/nnsCz100.bz
```

To navigate through the boot parameters, press Enter and the next parameter appears on the following line.

You can navigate through the entire list this way. To go back to a previous line, type a hyphen (-) and press Enter. Any value that you enter entirely overwrites the existing value and does not append to it.

3. To change a boot parameter, type the new value that you want to use next to the old value. For example, if you want to change the image you are using, type the new filename next to the old one. You can clear the contents of a parameter by typing a period and then pressing Enter.

```
ORACLE(configure)# bootparam  
'.' = clear field; '-' = go to previous field; q = quit  
  
Boot File      : /boot/nnsCz100.bz /boot/nnsCz200.bz
```

When you have scrolled through all of the boot parameters, the system prompt for the `configure terminal` branch displays.

NOTE: These changed parameters will not go into effect until reboot. Also, be aware that some boot parameters may also be changed through PHY and Network Interface Configurations.

```
ORACLE(configure)#
```

4. Exit the `configure terminal` branch.
5. Reboot the system for the changes to take effect.

The ACLI **reboot** and **reboot force** commands initiate a reboot. With the **reboot** command, you must confirm that you want to reboot. With the **reboot force** command, you do not have to make this confirmation.

```
ORACLE# reboot force
```

The system completes the full booting sequence. If necessary, you can stop the auto-boot at countdown to fix any boot parameters.

If you configured boot parameters correctly, the system prompt displays and you can go ahead with configuration, management, or monitoring tasks.

 **Note:**

If you configured the boot parameters incorrectly, the system goes into a booting loop and displays an error message. Press the space bar to stop the loop. Correct the error in the boot parameter, and reboot the system.

Change Boot Parameters by Interrupting a Boot in Progress

To access and change boot parameters by interrupting a boot in progress:

1. When the system is in the process of booting, you can press the space bar on your keyboard to interrupt when you see the following message appear:

```
Press the space bar to stop auto-boot...
```

2. After you stop the booting process, you can enter the letter **p** to display the current parameters, the letter **c** to change the boot parameters or the **@** (at-sign) to continue booting.

```
[Boot]: c
'.' = clear field; '-' = go to previous field; q = quit
Boot File   : /boot/nnScz100.bz
```

To navigate through the boot parameters, press Enter and the next parameter appears on the following line.

You can navigate through the entire list this way. To go back to a previous line, type a hyphen (-) and press Enter. Any value that you enter entirely overwrites the existing value and does not append to it.

3. To change a boot parameter, type the new value that you want to use next to the old value. For example, if you want to change the image you are using, type the new filename next to the old one.

```
'.' = clear field; '-' = go to previous field; q = quit
Boot File   : /boot/nnSCz100.bz /boot/nnSCz200.bz
```

4. After you have scrolled through the complete list of boot parameters, you return to the boot prompt. To reboot with your changes taking effect, type @ (the at-sign), and press Enter.

```
[Acme Packet Boot]: @
```

The system completes the full booting sequence, unless there is an error in the boot parameters.

If you have configured boot parameters correctly, the system prompt displays and you can go ahead with configuration, management, or monitoring tasks.

 **Note:**

If you have configured the boot parameters incorrectly, the system goes into a booting loop and displays an error message. Press the space bar to stop the loop. Correct the error, and reboot your system.

7

Formatting the Disk Volume

After the system has booted the Oracle Communications product for the first time, the hard disk must be formatted from the application. Remember that, at the chassis level, the disk system must be presented as a single logical drive.

 **Note:**

For systems without a separate boot flash, the **format system-disk** command combines the **/opt** and **/opt/crash** into a single system partition. This includes all systems other than the Acme Packet 4500, 4600, 6100 and 6300.

40GB or Less Format Plan

When formatting a 40 GB or smaller storage device, no data partitions are created. The default partitioning scheme is as follows:

Table 7-1 System Default Format Plan (40 GB max):

Location	Volume Name	Volume Size
system partition	/opt	8 GB
system partition	/opt/crash	32 GB

40GB or More Default Format Plan

When formatting a storage device larger than 40 GB, /mnt/sys and /mnt/app volumes are created in the data partition. Their relative sizes are based on the drive's size.

Table 7-2 System Format Plan (40 GB +):

Volume Number	Volume Name	Volume Size
system partition	/opt	8 GB
system partition	/opt/crash	2 x RAM size (not less than 8 GB)
data partition	/mnt/sys	20% remaining space
data partition	/mnt/app	80% remaining space

40GB or More Custom Format Plan

You can customize the format plan when a storage device larger than 40 GB is installed in your system. Before formatting the storage device, plan the number of

volumes, volume names, and relative percentage of storage device disk space. A maximum of 4 volumes in the data partition are allowed.

Table 7-3 Custom System Format Plan (40 GB +):

Volume Number	Volume Name	Volume Size
systempartition	/opt	8 GB
system partition	/opt/crash	2 x RAM size (not less than 8 GB)
data partitions 1-4	/mnt/<user-label>	user-defined percentage of remaining space

Formatting Procedure

The **format** command requires one of the following arguments:

- **system-disk** — formats and creates the 2 system partitions: /opt and /opt/crash
- **data-disk** — formats and creates 1 or more data partitions with the default (/mnt/sys and /mnt/app) or user-defined volumes
- **hard-disk** — formats and creates both the system partition and data partition

After the drive(s) are formatted, the system mounts the newly created partitions.

Use the **format hard-disk** command and argument after first system start up to fully initialize the media. The command's other arguments are applicable any other time.

Note:

The format command may only be executed if certain tasks like local CDR and HDR generation are not active. Remove any boot time configuration for these features and reboot the system before attempting to format the hard-disk. In addition, ensure that your device is not passing traffic while you format any partition.

The following example shows the format command process.

```
ORACLE# format hard-disk
WARNING: Please ensure device is not currently in use by any
applications before proceeding
Continue [y/n]?: y
The following system partitions will now be created:
1: /opt          8000000 bytes
2: /crash       16218284032 bytes
Create the system partitions and filesystems as configured above [y/
n]?: y
*****
WARNING: All system logs and data on the disk will be
permanently erased and unrecoverable.
Are you sure [y/n]?: y
The format process will take a few minutes. Once
```

```
the format process begins, it cannot be stopped.
Please do not power down or reboot the system until
the format process is complete.
Continue [y/n]?: y
Suspending logging to hard disk
Stopping tLogCleaner task
Relocating logging onto RAM drive
Initializing /opt/ Cleaner
Starting tLogCleaner task
*** Removing previous system partitions - please wait ***
*** Creating new system partitions - please wait ***
*** Formatting partition /opt. Please wait... ***
[...]
This filesystem will be automatically checked every 23 mounts or
180 days, whichever comes first. Use tune2fs -c or -i to override.
*** Formatting completed successfully ***
*** Formatting partition /crash. Please wait... ***
[...]
This filesystem will be automatically checked every 31 mounts or
180 days, whichever comes first. Use tune2fs -c or -i to override.
*** Formatting completed successfully ***
e2fsck 1.41.14 (22-Dec-2010)
opt: clean, 11/1960 files, 1323/7812 blocks
e2fsck 1.41.14 (22-Dec-2010)
crash: clean, 11/991232 files, 104681/3959542 blocks
```

This section of the format hard-drive walk-through shows the data partition creation. The following system output shows that the user has chosen to define a custom data partition scheme by typing "n" at the following prompt.

```
Use factory default data partitions [y/n]?:n
```

In this case, the user creates three partitions.

```
Suspending logging to RAM drive
Stopping tLogCleaner task
Relocating logging onto hard disk
Initializing /opt/ Cleaner
Starting tLogCleaner task
Disk space used by system:
    16226317824 bytes
Use factory default data partitions [y/n]?: n
Enter the number of data partitions to create: 3
Total unallocated space = 100 %
Enter the name of volume 1 (or 'q' to quit): VOLUME1
Enter the size of the volume (in %): 20
Total unallocated space = 80 %
Enter the name of volume 2 (or 'q' to quit): VOLUME2
Enter the size of the volume (in %): 40
Total unallocated space = 40 %
Enter the name of volume 3 (or 'q' to quit): VOLUME3
Enter the size of the volume (in %): 40
The following data partitions will now be created:
/VOLUME1 96776308838 bytes
```

```
/VOLUME2 193552617676 bytes
/VOLUME3 193552617676 bytes
Create the data partitions and filesystems as configured above [y/n]?: y
*****
WARNING: All non-system data on the disk will be
permanently erased and unrecoverable.
Are you sure [y/n]?: y
The format process will take a few minutes. Once
the format process begins, it cannot be stopped.
Please do not power down or reboot the system until
the format process is complete.
Continue [y/n]?: y
*** Beginning format process ***
*** Removing previous data partitions - please wait ***
*** Creating new data partitions - please wait ***
*** Formatting partition /VOLUME1. Please wait... ***
mke2fs 1.41.14 (22-Dec-2010)
[...]
Writing inode tables: done
Creating journal (32768 blocks): done
Writing superblocks and filesystem accounting information: done
This filesystem will be automatically checked every 37 mounts or
180 days, whichever comes first. Use tune2fs -c or -i to override.
*** Formatting completed successfully ***
*** Formatting partition /VOLUME2. Please wait... ***
mke2fs 1.41.14 (22-Dec-2010)
[...]
Writing inode tables: done
Creating journal (32768 blocks): done
Writing superblocks and filesystem accounting information: done
This filesystem will be automatically checked every 23 mounts or
180 days, whichever comes first. Use tune2fs -c or -i to override.
*** Formatting completed successfully ***
*** Formatting partition /VOLUME3. Please wait... ***
mke2fs 1.41.14 (22-Dec-2010)
[...]
Writing inode tables: done
Creating journal (32768 blocks): done
Writing superblocks and filesystem accounting information: done
This filesystem will be automatically checked every 31 mounts or
180 days, whichever comes first. Use tune2fs -c or -i to override.
*** Formatting completed successfully ***
*** Format finished successfully
New partitions have been created ***
*** Mounting partitions ***
e2fsck 1.41.14 (22-Dec-2010)
VOLUME1: clean, 11/5914624 files, 418265/23626953 blocks
/VOLUME1 mounted
e2fsck 1.41.14 (22-Dec-2010)
VOLUME2: clean, 11/11821056 files, 789884/47254150 blocks
/VOLUME2 mounted
e2fsck 1.41.14 (22-Dec-2010)
VOLUME3: clean, 11/11821056 files, 789884/47253628 blocks
/VOLUME3 mounted
```

Power cycle the system after format is complete. You can re-enable any tasks that may have conflicted with the format, including local CDR and HDR generation, after the power cycle is complete.

 **Note:**

If you find that, after the first reboot, the system has not created new partitions, perform another reboot to resolve this issue, which is created by an incorrect dynamic partition table refresh.

8

Interface Considerations for VM Platforms

Oracle Communications Session Delivery Product management and media traffic use separate network interfaces to connect to networks. The user configures these interfaces' properties, and then adapts other configuration elements to those interfaces. This configuration requires an understanding of the platform's mechanisms for using interfaces as well as the properties they can configure, as presented in this section.

Software Adaptations for Physical Interfaces

Having been installed on a virtual machine, Oracle Communications Session Delivery Product physical interface configuration elements provide for the applicable properties and same statistics measurements as are available for Acme Packet platform interfaces. With the exception of the MACTAB file, the abstraction of these elements from those physical ports is transparent to the user.

The user must independently understand the bus or virtual network paths used by their hardware to access physical ports. Refer to your hardware or virtual machine documentation for information about proper configuration and operation related to, for example, physical port location and speed. The user must consider these system and physical details when configuring Oracle Communications Session Delivery Product **phy-interface** elements, which is documented in the *ACLI Configuration Guide*.

Management vs. Media Interfaces

The template Oracle provides for VM deployment selects the correct interface detail for both management and media interfaces during deployment. If you are using a manual deployment approach or you wish to change interface selection after deployment, follow these guidelines for management and media interface selection:

- Always use emulated NICs for management interfaces
- Use either Paravirtual, PCI-pt or SR-IOV I/O mode for media interfaces

COTS Interfaces

Physical interfaces on COTS platforms are defined by the platform, with bussing and driver access fully independent of the Oracle Communications Session Delivery Product. Utilization of COTS interfaces by the Oracle Communications Session Delivery Product is the same as with VM interfaces. The Oracle Communications Session Delivery Product extracts physical interface MAC addressing from the platform and maps it to Oracle Communications Session Delivery Product configuration naming syntax to direct network traffic to physical interfaces based on the type of service they support.

COTS Network Interfaces

The Oracle Communications Session Delivery Product maps interfaces upon first startup, based on the hardware's NIC configuration. The software looks to configure 2 management and 2 service ports for the common 4 on-board NIC configuration. The presence of additional NIC cards maps an additional management interface to the on-board NICs, leaving only a single service interface on the on-board NIC. The Oracle Communications Session Delivery Product provides a means of testing and changing physical interface assignment.

The Oracle Communications Session Delivery Product instantiates the first management port, eth0/wancom0, when it boots. The boot parameters define this addressing to enable initial management access via IP for purposes including network boot. The user does not need to configure physical (or network) interfaces for wancom0.

The interface-mapping Branch

The **interface-mapping** branch resides at the root level of the Oracle Communications Session Delivery Product's ACLI. It contains a group of commands used for mapping physical interfaces on virtual machines or COTS platforms to the Oracle Communications Session Delivery Product application's configuration. The system accomplishes this using physical interface MAC addresses and ACLI configuration interface naming.

Sample default Oracle Communications Session Delivery Product interface mapping is presented below, using the **interface-mapping, show** command:

```
ORACLE(interface mapping)# show
Interface Mapping Info-----
Eth-IF  MAC-Addr                Label
wancom0 00:50:88:BC:11:12        #generic
wancom1 00:50:88:BC:61:6C        #generic
wancom2 00:50:88:BC:11:C7        #generic
spare   00:50:88:BC:61:12        #generic
s0p0    00:50:88:BC:71:79        #generic
s1p0    00:50:88:BC:21:FF        #generic
s0p1    00:50:88:BC:41:A2        #generic
s1p1    00:50:88:BC:31:AC        #generic
s0p2    FF:FF:FF:FF:FF:FF        #dummy
s1p2    FF:FF:FF:FF:FF:FF        #dummy
s0p3    FF:FF:FF:FF:FF:FF        #dummy
s1p3    FF:FF:FF:FF:FF:FF        #dummy
```

You can check or change **phy-interface** to MAC address configurations using the names shown under the Eth_IF column. You can identify the two types of physical interfaces that apply to the Oracle Communications Session Delivery Product, by the naming conventions:

- Management interfaces, shown above as wancom0, wancom1 and wancom2
- Media interfaces, shown above as s0p0, s0p1, s1p0 and s1p1

It is recommended that the user configure physical interfaces using the naming in the **Eth-IF** column above on COTS and VM platforms. These conventions, which simply use 's' for slot and 'p' for port, are visible in the **interface-mapping, show** output.

The default interface mapping assignment assumes four interfaces on the VM. If deployed with less than four, the user may need to re-assign the resulting interface mapping, which they can verify using the **interface-mapping, show** command after system start-up. If the mapping is wrong, the **interface-mapping, show** command allows the user to correct it. The most likely change would be to swap the wancom1 mapping with a viable media interface.

Working with the interface-mapping branch

Interface identification on the Oracle Communications Session Delivery Product is based on a system-level file called MACTAB that maps interface MAC addresses to interface naming that can be applied within Oracle Communications Session Delivery Product configuration. In most cases, users retain the default mapping. The **interface-mapping, show** command provide access to commands that allow the user see, change and even locate each interface based on this mapping. The MACTAB file is stored in the `/boot` folder. The MACTAB file ensures that interface mapping is persistent, and therefore usable, by your configuration regardless of changes to the system.

The **interface-mapping, show** command displays the current mapping. An example of a MACTAB file that a user has configured is provided below.

```
ORACLE(interface-mapping)#show
Interface Mapping Info
=====
Eth-IF      MAC-Addr          Label
wancom0     00:16:3E:30:00:2A # ctrl port, onboard MAC
wancom1     00:16:3E:30:00:2B # 2nd ctrl port, onboard MAC
s0p0        00:16:3E:30:00:2C # First media interface
s1p0        00:16:3E:30:00:2D # Second media interrface
=====
```

interface-mapping

The following table lists the **interface-mapping** commands along with their descriptions.

Command	Description
<code>interface-mapping show</code>	Display the existing content of <code>/boot/mactab</code> file, with the mapping information of all the available Ethernet Interface Names versus Physical Interface MAC addresses, along with any customer provided label information.
<code>interface-mapping locate <ethernet if name> <seconds></code>	Lets you visually locate the Ethernet media interface. One way to achieve this is to flash the LED of the physical interface when its device name is located. This parameter indicates, in seconds, when the flashing occurs on the LED.

Command	Description
interface-mapping label <ethernet if name> labeling text	Lets you label the Ethernet interface identified by <eth-if-name> with a text string you define. For example, you can use a label that is meaningful for your network layout. This label is stored and then displayed as # string after the MAC address for the Ethernet interface in the /boot/mactab file.
interface-mapping delete <ethernet if name>	Delete an unused Ethernet interface. The unused Ethernet interface could be result of changing network configuration. For example, if you replace an old NIC with a new one, the system writes the new one into mactab file, but does not delete the old one. A confirmation step appears with warning message. When you confirm the action, this entry is deleted from /boot/mactab file.
interface-mapping swap <ethernet if name1> <ethernet if name2>	Swap the mapping of Ethernet interface names against the available MAC physical interfaces. For example, you can first execute the interface-mapping show command to display the current information. interface-mapping show wancom0 00:16:3E:30:00:2A # control port, onboard MAC wancom1 00:16:3E:30:00:2B # 2nd control port, onboard MAC s0p0 00:16:3E:30:00:2C # PCI left side s1p0 00:16:3E:30:00:2D # PCI right side Then you can execute the interface-mapping swap command. interface-mapping swap s0p0 s1p0 wancom0 00:16:3E:30:00:2A # control port, onboard MAC wancom1 00:16:3E:30:00:2B # 2nd control port, onboard MAC s0p0 00:16:3E:30:00:2D # PCI right side s1p0 00:16:3E:30:00:2C # PCI left side A warning message appears. Once you confirm the action, the MAC addresses and their corresponding labels are swapped in the /boot/mactab/file.

 **Note:**

The **delete** and **swap** commands require a reboot to activate the new MACTAB.

Serial Interfaces

In lieu of IP management access, serial access provides the user with direct access to the Oracle Communications Session Delivery Product ACLI. The user must identify how their system allows for serial access. The serial interface can be a critical component of VM interface configuration as the user can make MACTAB changes via the serial interface without interrupting their own access during that management procedure.

Access to the Oracle Communications Session Delivery Product serial interface is dependent on platform. Familiarity with the platform is required to understand serial configuration.

Virtual machine management software provides a simulated serial port service from which the user gets direct serial access to each system. See your virtual machine manager's documentation for instructions on accessing these serial ports.

Serial port configuration, via boot parameters, is the same across all platforms.

9

Flash Drive Installation via Boot Media Creator

The Boot Media Creator (BMC), is provided within Oracle Session Delivery software distributions. BMC is an application that allows the user to create a USB drive from which they can install software.

Users identify the distribution that includes BMC by filename. These filenames start with the software version, appended with the text `-img-usb.exe`.

For example, `nnSCZ725.64-img-usb.exe` is the BMC distribution of software version S-CZ7.2.5.

Run this executable on a Windows systems to create USB installation media. The provisioned USB flash drive will then be used for software installation via Netra/ COTS USB interfaces.


Creating a Build Image

Use the Boot Media Creator (BMC) to write a bootable software image to a USB flashdrive.

Use the following procedure to create a USB stick containing a build image.


1. Download the BMC version you need to use. This filename begins with the software version and appended with the suffix **-img-usb.exe**.
2. Insert your USB stick. Note that BMC erases all data from this stick during the creation procedure.
3. Start the BMC application. The image below displays BMC version 1.3.2.1. Your version may differ, but the user procedure is intuitive and similar across versions.

Figure 9-1 BMC - Step 3




4. Click **Next**. Select the embedded image.

Figure 9-2 BMC - Step 4



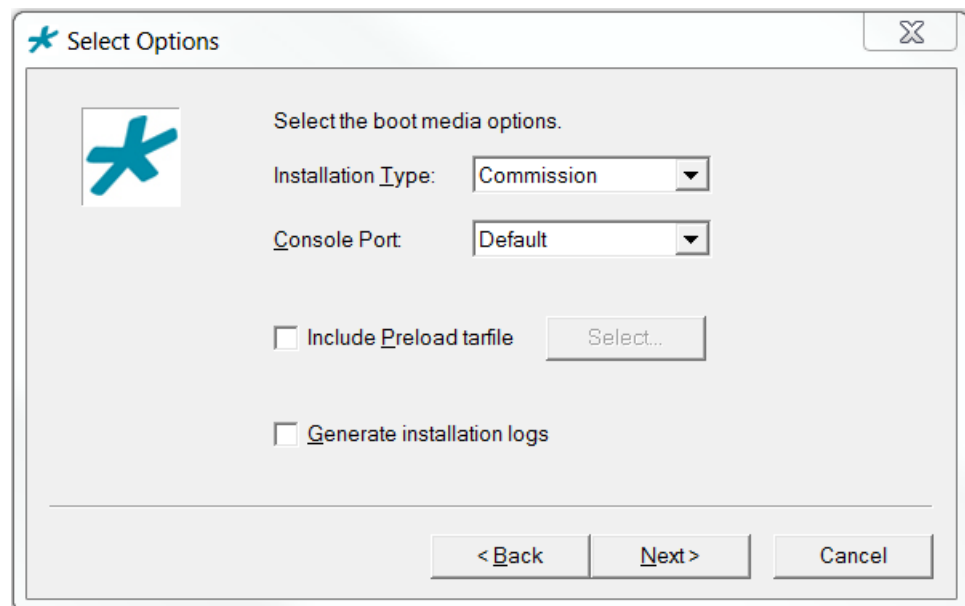
5. Click **Next**. BMC detects and displays your USB stick. You can insert and reinsert your stick, if needed.

Figure 9-3 BMC - Step 5



6. Highlight your stick and click **Next**. BMC displays a dialog allowing multiple options.

Figure 9-4 BMC - Step 6

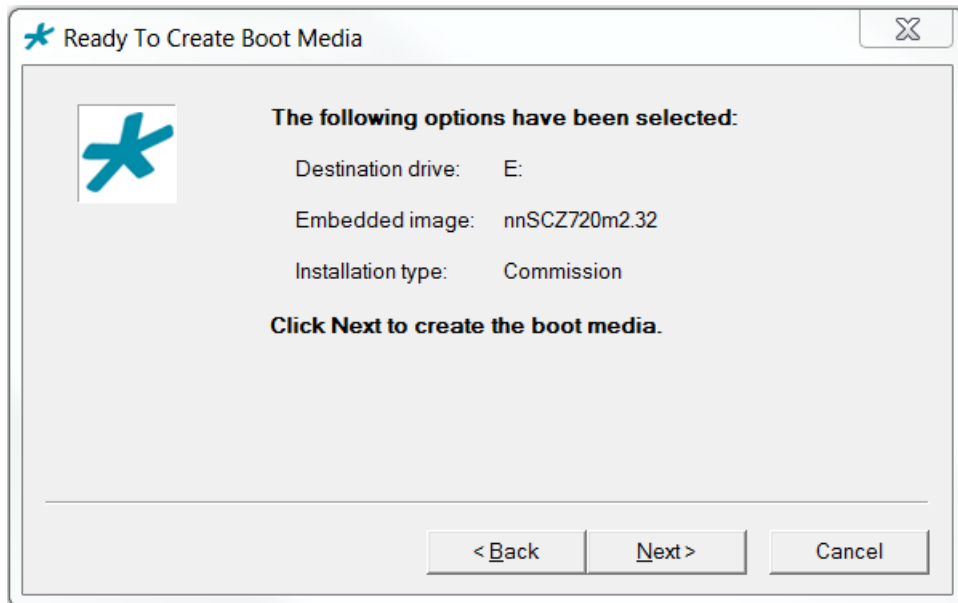


Note that the **Include Preload tarfile** and **Generate installation logs** options are typically used in a manufacturing or staging environment. User settings should be:

- Installation Type: Commission
- Console Port: Default

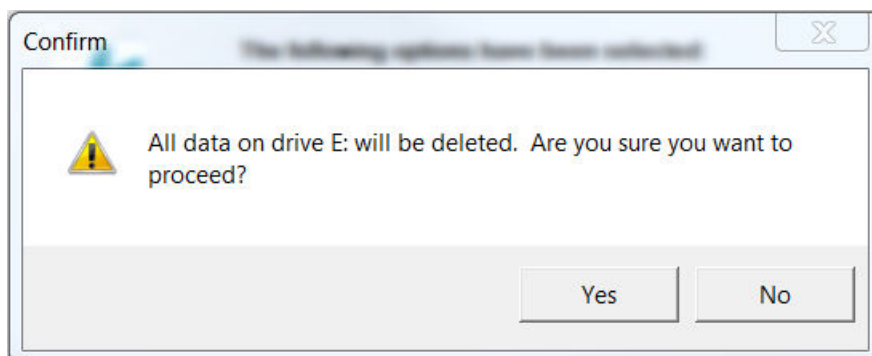
- Include Preload tarfile: unchecked
 - Generate installation logs: unchecked
7. Click **Next**. BMC confirms your settings.

Figure 9-5 BMC - Step 7




8. Click **Next**. BMC confirms that it will delete all data on your USB.

Figure 9-6 BMC - Step 8




9. Click **Yes**. BMC writes to the USB stick and indicates when it is finished.

Figure 9-7 BMC - Step 9



10. Click **Next** after the write operation is complete.

Figure 9-8 BMC - Step 10



11. Click **Back** to make another copy, or **Finish** to exit BMC.

Remove the USB flash drive when complete to ensure that the computer does not attempt to install your software during the next boot cycle.

10

Software Upgrade

This section provides information about how to upgrade your Oracle Communications Session Delivery Product software image. Note that proprietary Acme Packet hardware is normally delivered with an operational image ready to run. Your deployment, however, may require a different version than what was delivered, making software upgrade a required step in platform preparation.

Upgrade Checklist

Before upgrading the Oracle Communications Session Delivery Product software:

1. Obtain the name and location of the target software image file from either Oracle Software Delivery Cloud, <https://edelivery.oracle.com/>, or My Oracle Support, <https://support.oracle.com>, as applicable.
2. Provision platforms with the 64-bit Oracle Communications Session Delivery Product image file in the boot parameters. 64-bit image files are recognized by the "64" between the image revision and file extension. e.g., nnSCZ720.64.bz .
3. Verify the integrity of your configuration using the ACLI **verify-config** command.
4. Back up a well-working configuration. Name the file descriptively so you can fall back to this configuration easily.
5. Refer to the Oracle Communications Session Delivery Product Release Notes for any caveats involving software upgrades.

Download the Software

To get the Oracle Communications Session Delivery Product software, go to the Oracle Cloud Software Delivery website. With an account and a product license, you can download the software.

Before You Begin

Confirm that you have an account with Oracle and a license for the product that you want to download.

Download the Media Pack, software, and License Document for the product and platform that you want. In step 8, Oracle recommends that you view the Readme before attempting the software download.

1. Go to Oracle Cloud Software Delivery at <https://edelivery.oracle.com/>, and sign in.
2. On the Terms and Restrictions page, in the Oracle Trial License Agreement section, select one of the following:
 - If you want a trial license, select **Yes**.
 - If you have a license, select **Or**.
3. On the Terms and Restrictions page, in the Export Restrictions section, select **Yes**.
4. Click **Continue**.

5. On the Media Pack Search page, do the following:
 - Select a Product Pack
 - Select a Platform
6. Click **Go**.

The system displays a list of software for the selected product pack and platform.
7. Select the product that you want to download, and click **Continue**.
8. On the Download page, do the following:
 - a. Click **Readme** for download instructions and information about the Media Pack.
 - b. Click **Download** for the product download.
 - c. Click **Download** for the License Document download.
9. On the Oracle Software Delivery Cloud tool bar, click **Sign Out**.

Next Steps

You must unzip all of the files associated with a specific product into the same directory. You must also keep the directories for different products separate from each other. The directory in which you unzip the product files will be the staging area from where you will install the software.

Stand-alone Upgrade

This process incurs system downtime; your Oracle Communications Session Delivery Product stops passing traffic for a period of time. Please plan for your standalone upgrade accordingly.

Check /boot for free space

- On the OC-SDP, check for adequate space in the `/boot` volume to upload the new boot image and bootloader. Use the **show space boot** command.

```
OC-SDP# show space boot
boot: 24759488/25760512 bytes (99%) remaining
```

You may delete files from an SFTP client if you need to free space.

The command **check-space-remaining boot** performs the same function as **show space boot**.

Upload the Stage 3 Boot Loader and System Image

Whenever you upgrade the software image, upload the stage 3 boot loader and the new system image file to the system. The stage 3 boot loader is backward compatible with application software.


The Stage3 boot loader is generally backward compatible with previous releases, but Oracle recommends that the Stage3 boot loader be installed from the same Major.Minor version as the system image. It is not normally necessary to update the boot loader when installing a maintenance or patch release when the Major.Minor release is the same.

System upgrades typically consist of transferring the new system image and stage3 boot loader to the system and setting boot parameters to the new system software. To ensure compatibility, copy the stage 3 boot loader to `/boot/bootloader` before you update the boot parameters to use the new software image file. The boot loader file must be renamed to `/boot/bootloader` on the target system with no file extension. When upgrading an HA pair, you must perform the upgrade procedure on each HA node.

Follow the steps below to upload the Stage3 boot loader and system image.

1. Obtain the stage 3 boot loader image file (*.boot).
2. Upload the stage 3 boot loader image file (*.boot) as `/boot/bootloader` to your system using an SSH File Transfer Protocol (SFTP) client.
3. Upload the new system software image (*.bz) to `/boot/`.
4. Validate the bootloader by rebooting the Oracle Communications Session Delivery Product after renaming the bootloader.

```
[Downloads]$ ls -la
total 148820
drwxr-xr-x  2 bob src      4096 Jun 17 15:16 .
drwxr-xr-x 28 bob src      4096 May 21 14:17 ..
-rw-r--r--  1 bob src 10164527 Jun 17 15:15 nnSCZ720.64.boot
-rw-r--r--  1 bob src 73849839 Jun 17 15:15 nnSCZ720.64.bz
[Downloads]$ sftp admin@172.30.46.20
admin@172.30.46.20's password:
Connected to 172.30.46.20.
sftp> cd /boot
sftp> put nnSCZ720.64.boot
Uploading nnSCZ720.64.boot to /boot/nnSCZ720.64.boot
nnSCZ720.64.boot                100% 9926KB   9.7MB/s
00:01
sftp> rm /boot/bootloader
sftp> rename nnSCZ720.64.boot /boot/bootloader
sftp> put nnSCZ720.64.bz
Uploading nnSCZ720.64.bz to /boot/nnSCZ720.64.bz
nnSCZ720.64.bz                   100%   70MB  14.1MB/s
00:05
sftp> bye
Received disconnect from 172.30.46.20: 11: Logged out.
[Downloads]$
```

 **Note:**

Strictly speaking, the Stage3 boot loader is ready for operation after upload and filename change, but validating it before booting the new system software is good practice.

Software Upgrade Procedure

The following procedure describes how to upgrade a Oracle Communications Session Delivery Product with a new software image. In this procedure, the image file is

located on the Oracle Communications Session Delivery Product's local file system in `/boot`.

To upgrade a software image on a stand-alone system:

1. Change the boot configuration parameters to use the new image.

In the ACLI configure terminal menu, type **bootparam** and press <Enter> to display the list of boot parameters. Scroll through the boot parameters by pressing Enter. Stop when you reach the **file name** boot parameter and type the appropriate file name next to the previous file name. Press <Enter> to continue scrolling through the boot parameters.

The following example uses the filenames `/boot/nnSCZ720m5.64.bz` and `/boot/nnSCZ730.64.bz`.

```
OC-SDP1# configure terminal
OC-SDP1(configure)# bootparam
'.' = clear field; '-' = go to previous field; ^D = quit
boot device           : eth0
processor number      : 0
host name             : boothost
file name             : /boot/nnSCZ720m5.64.bz /boot/nnSCZ730.64.bz
```

2. Reboot the Oracle Communications Session Delivery Product using the **reboot** command.

The Oracle Communications Session Delivery Product should now be successfully running the new release.


HA Upgrade

In the descriptions and processes outlined below, ORACLE-1 is initially the active system and ORACLE-2 is initially the standby system. Please read the following procedures carefully before beginning the upgrade. If necessary, you can back out of the upgrade once during the upgrade procedure and once after you have completed the upgrade procedure.



Note:

See the diagram below, which addresses how the HA Upgrade procedures' sequence includes rebooting the standby system first to avoid service interruption.

Figure 10-1 Configure and Reboot Standby First for HA Procedures

HA Upgrade Procedure

This procedure upgrades HA deployments.

Note:

In the procedure below, ORACLE-1 is the active system and ORACLE-2 is the standby system. The standby system should be rebooted first.

1. Confirm that ORACLE-1 and ORACLE-2 start up and are synchronized.

Ensure that the running and current configurations on ORACLE-1 and ORACLE-2 have the same number. In the following examples, all of the configuration versions are 5.

On ORACLE-1 and ORACLE-2, use the **show health** command to ensure all processes are synchronized.

On ORACLE-1, show the current configuration version by using the ACLI **display-current-cfg-version** command. Then use the same command on ORACLE-2 and be sure that its current configuration version is the same as the one on ORACLE-1.

```
ORACLE-1# display-current-cfg-version
Current configuration version is 5
ORACLE-1#
ORACLE-2# display-current-cfg-version
Current configuration version is 5
ORACLE-2#
```

On ORACLE-1, show the running configuration version by using the ACLI **display-running-cfg-version** command. Then use the same command on ORACLE-2

and be sure that its running configuration version is the same as the one on ORACLE-1.

```
ORACLE-1# display-running-cfg-version
Running configuration version is 5
ORACLE-1#
ORACLE-2# display-running-cfg-version
Running configuration version is 5
ORACLE-2#
```

2. On ORACLE-2, before loading the software image to the flash, check the remaining space in the `/boot` directory using the ACLI **show space boot** command.

```
ORACLE-2# show space boot
boot: 24759488/25760512 bytes (99%) remaining
ORACLE-2#
```

If you see less than 50% of the space remaining, delete older stored firmware images to make space.

At a minimum, we recommend that you leave the `diags.gz` file and the currently running release on the flash memory (in the event that a rollback is required).

3. Upload the OC-SDP software image file and stage three bootloader to the `/boot` directory using an SFTP client. (See the instructions on updating the Stage 3 Bootloader.)
4. Change the boot configuration parameters on ORACLE-2 to use the appropriate new release software image.

Note:

From the point that you upgrade the image file, do not make any configuration changes. Likewise, do not use the **save-config** or **activate-config** commands. Once you execute the **save-config** command, the configuration can not be guaranteed to be backward compatible should you have to back out of the upgrade. Access the boot parameters on ORACLE-2:

- In the ACLI configure terminal menu, type **bootparam** and press <Enter> to being displaying the list of boot parameters.

Scroll through the boot parameters by pressing <Enter>. Stop when you reach the file name boot parameter.

The following example uses the filenames `/boot/nnSCZ720m5.64.bz` and `/boot/nnSCZ730.64.bz`.

```
ORACLE-2# configure terminal
ORACLE-2(configure)# bootparam
'.' = clear field; '-' = go to previous field; ^D = quit
boot device      : eth0
processor number  : 0
```

```

host name          : boothost
file name          : /boot/nnSCZ720m5.64.bz /boot/
nnSCZ730.64.bz

```

As shown above, type the new Release file name next to the previous one, including the path. Press <Enter> to continue scrolling through the boot parameters.

Reboot ORACLE-2.

5. After ORACLE-2 has completed the boot process, use the **verify-config** command to confirm that the configuration has been upgraded properly.

```
ORACLE-2# verify-config
```

6. Confirm the ORACLE-2 is running the new boot image using the **show version** command.

```

ORACLE-2# show version
Acme Packet 4600 SCZ7.3.0
Build Date=09/09/15

```

7. Use the **show health** command to confirm that ORACLE-2 is the standby system.
8. As you did for ORACLE-2, upload the OC-SDP software image file and stage three bootloder to the /boot directory using an SFTP client. (See the instructions on updating the Stage 3 Bootloader.)
9. As you did for ORACLE-2, configure the boot parameters on ORACLE-1 to boot from the new software image. Then reboot ORACLE-1.

```
ORACLE-1# reboot
```

```
-----
WARNING: you are about to reboot this SD!
-----
```

```
Reboot this SD [y/n]?: y
```

Rebooting ORACLE-1 causes ORACLE-2 to become the active system in the HA node.

10. When ORACLE-1 is finished rebooting, use the **show health** command to confirm that it is in the standby state.

Note:

If you need to revert to the older image, use the HA Backout Procedure.

HA Backout Procedure

If you reach the point in your upgrade procedure where you have upgraded both Oracle Communications Session Delivery Products in the HA pair to a later release that you decide you no longer want to use, you can fall back to a previous release. This section shows you how to fall back to an older image with both systems in your HA node upgraded.

 **Note:**

In the procedure below, ORACLE-2 is the active system and ORACLE-1 is the standby system. The procedure uses these designations because the prior procedure results in ORACLE-2 as the active system. The standby system should be rebooted first.

To backout to a previous (older) release with the both OC-SDPs in the HA node upgraded:

1. Change the boot parameters on ORACLE-1 to use the appropriate Release S-CZ7.2.0 software image.

Using one of these methods, access the boot parameters on ORACLE-1:

- Reboot ORACLE-1 using any of the ACLI **reboot** commands. Stop the booting process by hitting the Space bar on your keyboard to halt boot-up when you see this message: Press any key to stop auto-boot.... Type a **c** and press Enter to begin displaying the boot parameters.
- In the ACLI configure terminal menu, type **bootparam** and press Enter to begin displaying the list of boot parameters.

Scroll through the boot parameters by pressing Enter. Stop when you reach the file name boot parameter.

The following example uses the filenames `/boot/nnSCZ730.64.bz` and `/boot/nnSCZ720.64.bz`.

```
ORACLE-1# configure terminal
OC-SDP1(configure)# bootparam
'.' = clear field; '-' = go to previous field; ^D = quit
boot device          : eth0
processor number     : 0
host name            : boothost
file name            : /boot/nnSCZ730.64.bz /boot/nnSCZ720.64.bz
```

In the example above, type the appropriate Release S-CZ7.2.0 file name next to the Release S-CZ7.3.0 file name. Press <Enter> to continue scrolling through the boot parameters.

Exit to the main Superuser prompt.

```
ORACLE-1(configure)# exit
```

2. Reboot ORACLE-1.
3. Using the ACLI **show version** command to confirm that you are using the appropriate release.

```
ORACLE-1# show version
Acme Packet 4600 SCZ7.2.0
Build Date=01/09/15
```

4. Initiate a switchover on ORACLE-1.

```
ORACLE-1# notify berpd force
```

At this point, ORACLE-1 becomes the active system running Release S-CZ7.2.0. ORACLE-2 is now the standby system running Release S-CZ7.3.0.

5. On ORACLE-2, change the boot parameters as you did in Step 1 of this procedure.
6. Reboot ORACLE-2.

A

Physical Interfaces on Acme Packet Platforms

Acme Packet platforms are prepared for operation, including having the software pre-installed, prior to shipment. Refer to each product's hardware installation document for physical chassis installation. The following sections provide technicians with physical interface information that is useful for post-installation cabling.

The Network Interface Units (NIUs) installed on Acme Packet proprietary hardware define the number of interfaces, hardware protocol, and connection speed available for media and signaling traffic. Global operational information on this information is presented immediately below. Platform-specific physical information is presented in the subsequent sections. Use this information to understand these interfaces at a high level and to physically identify them for verification and troubleshooting procedures.

There are two types of physical interfaces on Acme Packet hardware NIUs:

- Media interfaces are on the network interface unit (NIU); they are also referred to as network media ports.
- Management interfaces are also on the NIU; they are also referred to as network management ports.

The first management interface, referred to as wancom0 or eth0, handles device management traffic including:

- SNMP
- SSH
- SFTP
- ACP/XML
- Logs sent from the Oracle Communications Session Delivery Product
- The boot file used to boot the Oracle Communications Session Delivery Product from a remote file server

The wancom0 interface does not require that the user perform physical or network interface branch configuration procedures within the ACLI. Instead, users configure its address and mask in the platform's boot parameters. Note that wancom0 uses the boot parameter's default gateway setting unless the system-config's default gateway is configured.

Users configure the wancom1 and wancom2 management interfaces for high availability (HA) state replication. For HA, these interfaces are often directly connected by a crossover cable.

Media interfaces handle session signaling and/or session media traffic. Users must perform all media, wancom1 and wancom2 interface configuration at the ACLI.

The table below summarizes the physical interface configuration parameters, which interface they are applicable to, and whether they are required.

Parameter	Network Media Interface	Wancom1 and wancom2 Network Management Interfaces
name	R	R
operation-type	R	R
port	R	R
slot	R	R
virtual-mac	O	I
admin-state	R	I
auto-negotiation	R	I
duplex-mode	R	I
speed	R	I
wancom-health-score	I	O


R = Required, O = Optional, I = Invalid

Acme Packet 1100 Physical Interfaces

The Acme Packet 1100 back panel comprises its primary network interface unit (NIU). The back panel includes the majority of the Acme Packet 1100's external interfaces, including console, alarm, network management and media interfaces. There is an optional 1-port T1/E1 NIU available for sending signaling traffic to the WAN as a backup voice channel. There is also a USB port on the front panel.

The graphic below shows the location and labeling of the Acme Packet 1100 media and network management ports. This labeling is an important point of reference when you set up the **phy-interface** configuration element.

Figure A-1 Acme Packet 1100 - Rear View



The Acme Packet 1100 NIU includes the following ports (from left to right).

- USB—For use only by Oracle personnel.
- SER MGT (Console)—Provides serial access for administrative and maintenance purposes.
- NET MGT (wancom0)—The system uses these 10/100/1000 Base-T Ethernet ports for device management functions and HA.

- INT—The system uses this 10/100 Mbps port for signaling and media traffic that is outbound with respect to the Remote Office/Branch Office (ROBO). This traffic originates from the ROBO.
- EXT—The system uses this 10/100 Mbps port for signaling and media traffic that is inbound with respect to the Remote Office/Branch Office (ROBO). This traffic is being sent to the ROBO.
- UNUSED—This port is not currently operational.

The graphic also displays the T1/E1 NIU installed above the data interfaces. The T1/E1 interface supports an RJ48C connector using Shielded Twisted Pair cable.

The Acme Packet 1100 can be shipped pre-configured for basic operation. From a physical perspective, the INT and EXT ports are the same, but the configuration pre-defines a number of complex ACLI configurations that make the INT suitable for cabling to internal ROBO infrastructure and the EXT suitable for cabling to infrastructure external to the ROBO.

The table below lists the labeling of each interface on the NIU, as well as the applicable **operation-type** and **port** parameters in the **phy-interface** configuration element. Note that the slot parameter for network management ports is always zero (0). The operation-type parameter distinguishes between otherwise overlapping slot/port configuration.


NIU Label	Operation-type	Slot	Port
SER MGT	NA	0	0
Net Mgmt	Maintenance	0	1
INT	Media	0	0
EXT	Media	0	1

Acme Packet 3900 Physical Interfaces

The Acme Packet 3900 platform uses one Network Interface Unit (NIU) that contains all external interfaces with ports for T1 and E1, serial management, network management, USBs, and media management.

The following illustration shows the NIU labels and ports, which you need to know about when you perform the phy-interface configuration.

Figure A-2 Acme Packet 3900 - Rear View



Ports key

- T1/E1—For Time Division Multiplexing (TDM) quad span
- SER MGT—For console access for administrative and maintenance purposes
- MGMT0—For EMS control, RADIUS accounting, CLI management, SNMP queries and traps, and other network management functions
- MGMT1 and MGMT2—For High Availability (HA), or for network management with no HA configuration
- USB—For a storage device, or for installing software
- P0 - P3—For signaling and media traffic on copper or fiber optic cable

When performing the phy-interface configuration, refer to the following table for mapping each NIU label and operation-type to the appropriate slot and port parameters.

NIU Label	Operation-type	Slot	Port
Mgmt 0	Maintenance	0	0
Mgmt 1	Maintenance	0	1
Mgmt 2	Maintenance	0	2
P0	Media	0	0
P1	Media	0	1
P2	Media	1	0
P3	Media	1	1

Hardware Support

The Acme Packet 3900 hardware provides the following:

- 1 management interface at 1Gbps
- 4 media and signalling interfaces at 10/100/1000Mbps
- 1 HA interface at 10/100/1000Mbps
- 4 USB ports
- Hardware transcoding support for up to 5 Digital Signal Processor (DSP) modules
- 1 quad-span Time Division Multiplexing (TDM) PCIe card

SNMP Hardware Reporting

The Acme Packet 3900 platform relies on a specific set of MIB objects, in addition to the standard MIB objects.

The Acme Packet 3900 platform supports MIB objects for power supplies, fans, temperature sensors, system information, transcoding DSP(s), wancom ports, media ports, and the product OID. The Standard MIBs (such as MIB-2 objects) are supported.

The Acme Packet 3900 monitors the following environmental parameters by way of SNMP:

Updates to sysObjectID OID in the ap-products.mib.

- Updates the apNetNet 3000Series object to include the apNetNet 3900 object.

Updates to the entity OID in ap-entity-vendortype.mib.

- Updates the `apevPowerSupply` object to include the `apevPowerSupply 500 W` object.

Acme Packet 3900 MIBS Paths

Paths for Acme Packet 3900 MIBS.

SNMPv2-SMI::mib-2.47.1.1.1.1.2.1 = STRING: "Acme Packet 3900 Chassis"

SNMPv2-SMI::mib-2.47.1.1.1.1.2.2 = STRING: "Intel(R) Xeon(R) CPU D-1548 @ 2.00GHz"

SNMPv2-SMI::mib-2.47.1.1.1.1.2.3 = STRING: "495 Watt Power Supply"

SNMPv2-SMI::mib-2.47.1.1.1.1.2.4 = STRING: "500 Watt Power Supply"

SNMPv2-SMI::mib-2.47.1.1.1.1.2.5 = STRING: "Assy, 2-fan unit of 40x10"

SNMPv2-SMI::mib-2.47.1.1.1.1.2.6 = STRING: "Sensor of fan speed"

SNMPv2-SMI::mib-2.47.1.1.1.1.2.7 = STRING: "Assy, Acme Packet 3900 Main Board"

SNMPv2-SMI::mib-2.47.1.1.1.1.2.8 = STRING: "Sensor of temperature"

SNMPv2-SMI::mib-2.47.1.1.1.1.2.9 = STRING: "Management Port 0 10/100 Ethernet Copper"

SNMPv2-SMI::mib-2.47.1.1.1.1.2.10 = STRING: "Management Port 1 10/100 Ethernet Copper"

SNMPv2-SMI::mib-2.47.1.1.1.1.2.11 = STRING: "Management Port 2 10/100 Ethernet Copper"

SNMPv2-SMI::mib-2.47.1.1.1.1.2.12 = STRING: "Media port - Logical Slot 0 Port 0"

SNMPv2-SMI::mib-2.47.1.1.1.1.2.13 = STRING: "Media port - Logical Slot 0 Port 1"

SNMPv2-SMI::mib-2.47.1.1.1.1.2.14 = STRING: "Media port - Logical Slot 1 Port 0"


SNMPv2-SMI::mib-2.47.1.1.1.1.2.15 = STRING: "Media port - Logical Slot 1 Port 1"

Acme Packet 4600 Physical Interfaces

The Acme Packet 4600 supports a single network interface unit (NIU) that contains all external interfaces, including console, alarm, network management and media interfaces. There is currently one type of NIU available, which defines the supported cabling and speed.

The graphic below shows the Acme Packet 4600 NIU ports with labeling. This labeling is an important point of reference when you set up the **phy-interface** configuration element.

Figure A-3 Acme Packet 4600 - Rear View



The Acme Packet 4600 NIU includes the following ports (from left to right).

- Console—Provides serial access for administrative and maintenance purposes.
- Alarm—Dry contact alarm port.
- USB—The USB port is reserved for use by Oracle support employees only.
- Mgmt 0–Mgmt 2—The system uses these 10/100/1000 Base-T Ethernet ports are used for device management functions. The first interface, Mgmt 0, is for ssh access to the ACLI. The other two interfaces are used for state replication for High Availability (HA). For HA, connect these interfaces directly using a crossover cable.
- P4–P5—The system uses these 2 x 10GbE ports for signaling and media traffic.
- P0–P3—The system uses these 4 x GbE ports for signaling and media traffic.

⚠ WARNING:

Customers may use either the 2 x 10GbE ports or the 4 x GbE media ports. Using both P4-P5 and P0-P3 is not supported.

The table below lists the labeling of each interface on the NIU, as well as the applicable **operation-type** and **port** parameters in the **phy-interface** configuration element. The **slot** parameter for this platform is always set to 0. The operation-type parameter distinguishes between otherwise overlapping slot/port configuration.


NIU Label	Operation-type	Slot	Port
Mgmt 0	Maintenance	0	0
Mgmt 1	Maintenance	0	1
Mgmt 2	Maintenance	0	2
P0	Media	0	0
P1	Media	0	1
P2	Media	0	2
P3	Media	0	3

Acme Packet 6100 Physical Interfaces

The Acme Packet 6100 supports a single network interface unit (NIU) that contains all external interfaces, including console, alarm, network management and media interfaces. There is currently one type of NIU available, which defines the supported cabling and speed.

The graphic below shows the NIU front panel, which includes all ports and their labeling. This labeling is an important point of reference when you set up the **phy-interface** configuration element.

Figure A-4 Acme Packet 6100 - Rear View



The Acme Packet 6100 NIU includes the following ports (from left to right).

- Console—Provides serial access for administrative and maintenance purposes.
- Alarm—Dry contact alarm port.
- USB—For use only by Oracle personnel.
- Mgmt0 to Mgmt2—The system uses these 10/100/1000 Base-T Ethernet ports for device management functions. The first interface, Mgmt 0, is for ssh access to the ACLI. The other two interfaces are used for state replication for High Availability (HA). For HA, connect these interfaces directly using a crossover cable.
- SFP+ ports—The system uses these 2 x 10GbE ports for signaling and media traffic.

The table below lists the labeling of each interface on the NIU, as well as the applicable **operation-type** and **port** parameters in the **phy-interface** configuration element. Note that the media interfaces are not uniquely labeled with the chassis silkscreen. The table distinguishes between these using "left" and "right", with the perspective being the user looking at the NIU panel.

NIU Label	Operation-type	Slot	Port
Mgmt 0	Maintenance	0	0
Mgmt 1	Maintenance	0	1
Mgmt 2	Maintenance	0	2
USB	NA	NA	NA
NA (left)	Media	0	0
NA (right)	Media	0	1

Acme Packet 6300/6350 Physical Interfaces

The Acme Packet 6300/6350 integrates interfaces for console, alarm, USB, and network management into the rear of the chassis, above the power supplies. The system also includes three PHY slots.

Acme Packet 6300/6350 Management Interfaces

Management interfaces include 3 x 10/100/1000 Ethernet interfaces labeled Mgmt 0, 1, and 2. (These interfaces are often referred to as wancom0, 1, and 2.) Mgmt 0, is for SSH access to the ACLI. The other two interfaces are used for state replication

for High Availability (HA). For HA, connect these interfaces directly using a crossover cable.

The console port is serial and the alarm is "dry contact".

The USB port is reserved for use by Oracle personnel, only. The following illustration shows the ports.

Figure A-5 Acme Packet 6300/6350 Management Interfaces



The following table maps the reference numbers in the preceding illustration and the labeling on the management interfaces. It also, where applicable, lists the **operation-type**, **slot** and **port** parameters in the **phy-interface** configuration element for the management interfaces.

Graphic Label	NIU Label	Operation-type	Slot	Port
1	Console	NA	NA	NA
2	Alarm	NA	NA	NA
3	Mgmt 0	Maintenance	0	0
4	Mgmt 1	Maintenance	0	1
5	Mgmt 2	Maintenance	0	2
6	USB	NA	NA	NA

Acme Packet 6300

Acme Packet 6300 Media Interface


The Acme Packet 6300 contains three PHY card slots. Slots 0 and 1 support network-facing media interface cards. Slots 1 and 2 support transcoding cards. The Acme Packet 6300 PHY cards contain 2x10Gb Ethernet ports.

The Acme Packet 6300 only supports the 2x10GbE Network Interface Unit (NIU).

Note:

Do not put any 2x10Gb NIU into slot 2.

Figure A-6 Acme Packet 6300 Media Interfaces



The following table maps the reference numbers in the preceding illustration and the labeling on the media interfaces. It also lists the **operation-type**, **slot**, and **port** parameters in the **phy-interface** configuration element for the applicable interfaces. The table distinguishes between these using "left" and "right", with the perspective of you looking at the NIU panel.

Graphic Label	NIU Label	Operation-type	Slot	Port
1	NA (left)	Media	0	0
2	NA (right)	Media	0	1
3	NA (left)	Media	1	0
4	NA (right)	Media	1	1

Acme Packet 6350

Acme Packet 6350 Media Interface

The Acme Packet 6300 contains three PHY card slots. Slots 0 and 1 support network-facing media interface cards. Slots 1 and 2 support transcoding cards. The Acme Packet 6350 supports both the 2x10GbE NIU and 4x10GbE NIU. The Acme Packet 6350 can support more than one 2x10Gb Ethernet NIU, but it supports only one 4x10Gb Ethernet NIU.

If you use two 2x10GbE NIUs, put the first one in slot 0, the second one in slot 1. When using the 4x10GbE NIU, you must put it in slot 0.

Note:

Do not put the 2x10GbE NIU in slot 2, and do not put the 4x10GbE NIU in either slots 1 or 2.

The Acme Packet 6350 supports up to 4x10GbE ports each running at full duplex line rate for all packet sizes. The

following illustration shows the media ports numbering pattern.




Figure A-7 Acme Packet 6350 - Media Interface

The following table maps the reference numbers in the preceding illustration to the labeling on the media interfaces. It also lists the **operation-type**, **slot**, and **port** parameters in the **phy-interface** configuration element for the applicable interfaces. The table distinguishes between these using "left" and "right", with the perspective of you looking at the NIU panel.

Graphic Label	NIU Label	Operation-type	Slot	Port
0	NA (left)	Media	0	0
1	NA (left)	Media	0	1
2	NA (right)	Media	0	2
3	NA (right)	Media	0	3