

# Oracle® Enterprise Session Border Controller

## Installation and Platform Preparation Guide



Release E-CZ7.5.0

F19880-01

February 2020

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

ORACLE®

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

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 installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. 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 Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark 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	<b>Introduction to Platform Preparation and Software Deployment</b>	
	Acme Packet Engineered Platforms	1-1
	Virtual Machines	1-1
2	<b>The Oracle Communications ESBC as a Virtual Machine</b>	
	Hyper-threading and CPU Affinity	2-1
	Host Hypervisor CPU Pinning	2-1
	Configuration Overview	2-2
	VLAN Support	2-3
	Provisioning Entitlements	2-3
3	<b>Virtual Machine Platforms</b>	
	Create and Deploy on Oracle VM Manager	3-1
	Create and Deploy on KVM	3-3
	Create and Deploy on VMware®	3-5
4	<b>Boot Management</b>	
	Boot Loader Overview	4-1
	Considerations When Booting the Acme Packet 1100 and Acme Packet 3900	4-2
	Boot Parameters	4-2
	Boot Parameter Definitions	4-3
	Boot Flags	4-4
	Changing Boot Parameters	4-5
	Change Boot Parameters from the ACLI	4-5
	Change Boot Parameters by Interrupting a Boot in Progress	4-6
	Boot Management for the Acme Packet 4500	4-7
	Acme Packet 4500 Boot Loaders	4-8

Acme Packet 4500 Bootloader Prerequisites	4-9
Acme Packet 4500 Boot Parameters	4-9
Acme Packet 3820 and 4500 Boot Parameter Definitions	4-10
Acme Packet 3820 and 4500 Boot flags	4-11

## 5 Formatting the Disk Volume

---

40GB or Less Format Plan	5-1
40GB or More Default Format Plan	5-1
40GB or More Custom Format Plan	5-1
Formatting Procedure	5-2

## 6 Interface Considerations for VM Platforms

---

Software Adaptations for Physical Interfaces	6-1
The interface-mapping Branch	6-1
Working with the interface-mapping branch	6-2
interface-mapping	6-2
Serial Interfaces	6-4

## 7 Software Upgrade

---

Upgrade Checklist	7-1
Acme Packet 4500 Upgrade Requirements	7-1
Stand-alone Upgrade	7-1
Check /boot for free space	7-2
Upload the Stage 3 Boot Loader and System Image	7-2
Software Upgrade Procedure	7-3
HA Upgrade	7-3
HA Upgrade Procedure	7-4
HA Backout Procedure	7-6
HA Backout Procedure	7-7
Moving a Configuration	7-8
Backup Commands	7-8
Backing up the current configuration	7-9
Listing Backups	7-9
Copy the Backup to the destination	7-9
Restoring Backups	7-10

## 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-4
Acme Packet 4500 Physical Interfaces	A-5
Acme Packet 4600 Physical Interfaces	A-6
Acme Packet 6300/6350 Physical Interfaces	A-7

## B DPDK-Based OVS Installation

---

## List of Figures

---

3-1	Properties Dialog for OVF Deployment	3-6
3-2	OVF Deployment Status Dialog	3-7
A-1	Acme Packet 1100 - Rear View	A-2

## List of Tables

---

5-1	System Default Format Plan (40 GB max):	5-1
5-2	System Format Plan (40 GB +):	5-1
5-3	Custom System Format Plan (40 GB +):	5-2

# About This Guide

The *Installation and Platform Preparation Guide* explains the concepts and procedures necessary for platform preparation, power-on, and software deployment for the Oracle Communications Enterprise Session Border Controller (ESBC). The topics provide the information that you need after physical installation, and before you perform service configuration. The guide covers all supported platforms with information specific to each one. The information in the guide can also help you after product deployment, for example, to identify physical interfaces.

## Documentation Set

The E-CZ7.5.0 documentation set differs from previous releases with the addition of separate guides for installation, call traffic monitoring, and header manipulation rules. The content for the new guides was previously located in the *ACLI Configuration Guide*, which no longer contains such information. The documentation set also includes new guides for Federal Information Processing Standard (FIPS) compliance and the Admin Security licenses.

The following table describes the documents included in the ESBC E-CZ7.5.0 documentation set.

Document Name	Document Description
ACLI Configuration Guide	Contains conceptual and procedural information for configuring, administering, and troubleshooting the ESBC.
Administrative Security Guide	Contains conceptual and procedural information for supporting the Admin Security license, the Admin Security ACP license, and JITC on the ESBC.
Call Traffic Monitoring Guide	Contains conceptual and procedural information for configuration using the tools and protocols required to manage call traffic on the ESBC.
FIPS Compliance Guide	Contains conceptual and procedural information about FIPS compliance on the ESBC.
HMR Guide	Contains conceptual and procedural information for header manipulation. Includes rules, use cases, configuration, import, export, and examples.
Installation and Platform Preparation Guide	Contains conceptual and procedural information for system provisioning, software installations, and upgrades.
Release Notes	Contains information about the E-Cz7.5.0 release, including platform support, new features, caveats, known issues, and limitations.
Web GUI User Guide	Contains conceptual and procedural information for using the tools and features of the ESBC Web GUI.



## Related Documentation

The following table lists other documentation related to using the ESBC. You can find the listed documents on <http://docs.oracle.com/en/industries/communications/> in the "Session Border Controller Documentation" and "Acme Packet" libraries.

Document Name	Document Description
Accounting Guide	Contains information about the ESBC accounting support, including details about RADIUS accounting.
ACLI Reference Guide	Contains explanations of how to use the ACLI, as an alphabetical listings and descriptions of all ACLI commands and configuration parameters.
Acme Packet 4600 Hardware Installation Guide	Contains information about the components and installation of the Acme Packet 4600.
Acme Packet 6300 Hardware Installation Guide	Contains information about the components and installation of the Acme Packet 6300.
HDR Resource Guide	Contains information about the ESBC Historical Data Recording (HDR) feature. This guide includes HDR configuration and system-wide statistical information.
Maintenance and Troubleshooting Guide	Contains information about ESBC logs, performance announcements, system management, inventory management, upgrades, working with configurations, and managing backups and archives.
MIB Reference Guide	Contains information about Management Information Base (MIBs), Acme Packet'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.
Security Guide	Contains information about security considerations and best practices from a network and application security perspective for the ESBC family of products.

## Revision History

The following table describes the history of updates to this guide.

Date	Description
August, 2017	<ul style="list-style-type: none"> <li>Initial release</li> </ul>
October, 2017	<ul style="list-style-type: none"> <li>Adds content clarifying management vs media interface types for VNF</li> </ul>
February 2020	<ul style="list-style-type: none"> <li>Updates the Media Interface port numbers in "Acme Packet 1100 Physical Interfaces."</li> </ul>

# 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.

## 2

# The Oracle Communications ESBC as a Virtual Machine

You can deploy version E-CZ7.5.0 of the ESBC as a Virtual Machine (VM).

VM deployment types include:

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

Standalone ESBC VM deployment instances are always supported. Support within an orchestrated environment is dependent on orchestrator and ESBC 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 ESBC'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.

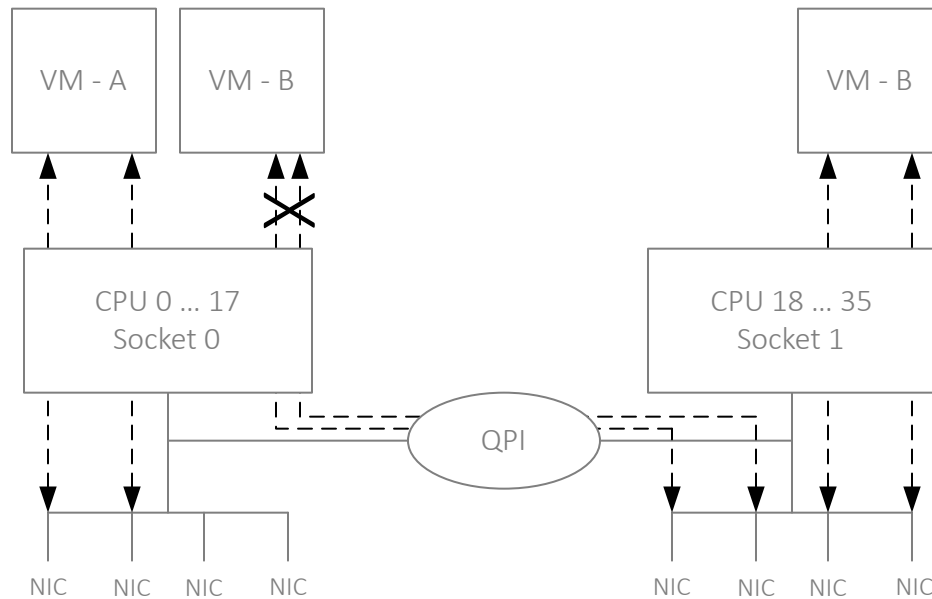
The user can learn how to configure CPU affinity via their hypervisor documentation.

## 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.



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

## Configuration Overview

Oracle Communications Enterprise Session Border Controller Virtual Machine (VM) deployments require configuration of the VM environment and, separately, configuration of the ESBC itself. VM-specific configuration on the ESBC 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 ESBC 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 ESBC 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 ESBC 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 ESBC-related configuration.

## VLAN Support

Version E-CZ7.5.0 supports VLANs within all supported virtual machine environments and interface architectures. Refer to the *ACLI Configuration Guide* for instructions on configuring VLANs on the ESBC. Note that when you configure a VLAN, the ESBC 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.

# 3

## Virtual Machine Platforms

Oracle distributes virtual machine templates, each containing a virtual disk image and sample configuration for the minimum 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/](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 ESBC 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 ESBC. 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=media1 \  
--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 ESBC 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 ESBC. 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 3-1 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 3-2 OVF Deployment Status Dialog**



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

# 4

## 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.

### Stage1 and Stage2 Boot Loaders

Stage1 and Stage2 are only used with the Acme Packet 4500 platform.

### 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 ESBC does not support uploading the boot loader by way of the Web GUI.

## Considerations When Booting the Acme Packet 1100 and Acme Packet 3900

The Acme Packet 1100 and Acme Packet 3900 use content in the **/boot/usbcPurpose** and **/boot/usdpMemory** files for operation. Do not change these files manually.

When booting, the Acme Packet 1100 and Acme Packet 3900 use the **/boot/usbcPurpose** and **/boot/usdpMemory** files when they are present, and creates them when they are not present. If you manually change these files, the Acme Packet 1100 and Acme Packet 3900 does not boot or function correctly.

If anyone made any changes to these files, delete the files before rebooting. The system re-creates the files correctly during the reboot.

 **Note:**

You might see the **/boot/cmdline file** on the system. The file is obsolete and does not present this problem.

## 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
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 Enterprise Session Border Controller 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 Enterprise Session Border Controllers 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.

### Note:

The EC2 platform does not support a writeable console, preventing the user from changing boot parameters by interrupting the system boot process.

Oracle recommends that you use management port 0 (wancom0) as the boot interface, and that your management network is either:

- directly a part of your LAN for management port 0
- accessible through management port 0

## 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; ^D = quit  
boot device          : eth0
```

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; ^D = quit
boot device      : eth0
processor number : 0
host name       : goose
file name       : /boot/nnPCz100.bz /boot/nnPCz200.bz
```

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

```
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 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.

```
Error loading file: errno = 0x226.
Can't load boot file!!
```

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.

```
[Acme Packet Boot]: c
'.' = clear field; '-' = go to previous field; ^D = quit
boot device      : wancom0
```

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; ^D = quit
boot device      : wancom0
processor number  : 0
host name        : goose
file name        : /code/nnPCz100.gz /code/nnPCz200.gz

```

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.

```
Error loading file: errno = 0x226.
Can't load boot file!!
```

Press the space bar to stop the loop. Correct the error, and reboot your system.

## Boot Management for the Acme Packet 4500

Boot management for the Acme Packet 4500 differs from other platforms and is, therefore, covered separately in this section. The following sections discuss boot management on the Acme Packet 4500 in detail. Boot Management includes the tasks needed to ensure the system is operating according to the users requirements after 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.

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.

## Acme Packet 4500 Boot Loaders

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 the user need not consider boot loader version during first install procedures. Application software upgrades do not update boot loaders. For this reason, the user needs to verify this compatibility manually. This section provides information about identifying the need and performing these updates.

Boot loaders handling on the Acme Packet 4500 differs from all other platforms. Important differences between the Acme Packet 4500 platform's boot loaders include:

- The Stage1 and Stage2 boot loaders reside in the **/boot** directory.
- The Stage1 boot loader is named **bootrom.sys** when it is deployed properly.
- The means of transferring boot loader files requires that the user execute system level commands after new file transfer is complete to perform the update.
- The **show version boot** command displays the Stage1 and Stage2 file creation date, which can be understood as the files' versions.
- When upgrading the Acme Packet 4500, make sure the stage2.gz file (or any file named stage2.gz) is not in the **/code** directory.

### Stage1 and Stage2 Boot Loaders

Stage1 and Stage2 boot loader upgrade is not required for every software upgrade. For this reason, these boot loaders are not provided along with every software distribution. Oracle provides version change information, when necessary, based on application release version.

### 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 nnECZ720.boot can be used with ECZ720, ECZ720m1, and so forth system software. But it should be upgraded when installing ECZ730 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.

## Acme Packet 4500 Bootloader Prerequisites

Customers planning to upgrade to Oracle Session Deliver software version 7 and beyond on Acme Packet 4500 hardware must first upgrade their boot loader to the image dated July 03, 2013 or later.

This boot loader upgrade includes applying the following two files:

- bootrom.gz
- stage2.gz

You must verify that your Acme Packet 4500 has been upgraded to the July 03, 2013 boot loader with the show version boot ACLI command.

```
ORACLE# show version boot
Bootloader Info
```

```
Stage 1: Jul 3 2013 13:16:30
Stage 2: Jul 3 2013 13:16:30
```

If the above information is not printed on your terminal, do NOT proceed with your upgrade. Contact TAC for instructions on how to continue.

### Note:

As you upgrade your boot loader in preparation to run version 7 and later software, NEVER place the stage2.gz file (or any file named stage2.gz) into the /code directory.

## Acme Packet 4500 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.
- Sets up the username and password for network booting from an external FTP server.

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.

The full set of Acme Packet 4500 boot parameters differs from those on other platforms, and appear as below:

```
ORACLE(configure)# bootparam
.' = clear field; '-' = go to previous field; ^D = quit
boot device           : eth0
processor number      : 0
host name             : acmepacket8
```

```

file name           : /boot/nnSC600.gz
inet on ethernet (e) : 10.0.1.57:ffff0000
inet on backplane (b) : 0.0.0.0
host inet (h)       : 10.0.1.5
gateway inet (g)    : 10.0.0.1
user (u)            : user
ftp password (pw)   : password
flags (f)           : 0x08
target name (tn)    : acmesystem
startup script (s)  : 0
other (o)           :

```

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

Notes:

- boot device—The boot device should be **eth0**.
- file name—The file name should start with **/boot/**.

## Acme Packet 3820 and 4500 Boot Parameter Definitions

The Acme Packet 3820 and 4500 series displays boot parameters when the user configures them after a boot interrupt. The system masks these additional parameters within 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 Acme Packet 3820 and 4500 boot parameters that are visible when the user configures after a boot interrupt.

Boot Parameter	Description
boot device	Management interface name and port number of the device from which an image is downloaded (e.g., wancom0 or eth0) from an external device.
processor number	Processor number on the backplane.
host name	Name of the boot host used when booting from an external device.
file name	Name of the image file to be booted; can be entered with the filename path. If you are booting from the flash memory, this filename must always match the filename that you designate when you FTP the image from the source to the system. When booting from internal flash memory, this filename must start with /boot); for example, /boot/nnECZ710.bz.
inet on ethernet (e)	Internet address of the system. This field can have an optional subnet mask in the form inet_adrs:subnet_mask. If DHCP is used to obtain the parameters, lease timing information may also be present. This information takes the form of lease_duration:lease_origin and is appended to the end of the field. In this parameter, the subnet mask ffff0000 = 255.255.0.0. When you use the ACLI acquire-config command, this is the IPv4 address of the system from which you will copy a configuration.
inet on backplane (b)	Internet address of the backplane interface, eth0. This parameter can have an optional subnet mask and/or lease timing information, such as e (inet on ethernet) does.

---

Boot Parameter	Description
host inet (h)	Internet address of the boot host used when booting from an external device.
gateway inet (g)	Internet address of the gateway to the boot host. Leave this parameter blank if the host is on the same network.
user (u)	FTP username on the boot host.
ftp password (pw)	FTP password for the FTP user on the boot host.
flags (f)	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 (tn)	Name of the system 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 systems in your network. This name can be 64 characters or less.
startup script (s)	For Oracle use only.
other (o)	For Oracle use only.

---

## Acme Packet 3820 and 4500 Boot flags

The following boot flags, configured in the boot parameter's **flags** parameter are for use on systems running on the Acme Packet 3800 and 4500 platforms.

- 0x00000008 Bootloader ~7 seconds countdown

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

- 0x00000002 Kernel debug enable
- 0x00000004 No auto boot
- 0x00000010 Debug sshd enable
- 0x00000020 Debug console enable
- 0x00001000 Userspace debug enable
- 0x00002000 acme.ko network module security override
- 0x00004000 Crashdump enable
- 0x00100000 Uniprocessor enable (SMP disable)
- 0x20000000 Fail-safe boot enable
- 0x40000000 Process startup disable (flatspin mode)



# 5

## 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 5-1 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 5-2 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 5-3 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.

# 6

## Interface Considerations for VM Platforms

Oracle Communications Enterprise Session Border Controller 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 Enterprise Session Border Controller 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 Enterprise Session Border Controller **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

### The interface-mapping Branch

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

Sample default Oracle Communications Enterprise Session Border Controller 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
slp0    00:50:88:BC:21:FF      #generic
s0p1    00:50:88:BC:41:A2      #generic
slp1    00:50:88:BC:31:AC      #generic
s0p2    FF:FF:FF:FF:FF:FF      #dummy
slp2    FF:FF:FF:FF:FF:FF      #dummy
s0p3    FF:FF:FF:FF:FF:FF      #dummy
slp3    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 Enterprise Session Border Controller, by the naming conventions:

- Management interfaces, shown above as wancom0, wancom1 and wancom2
- Media interfaces, shown above as s0p0, s0p1, slp0 and slp1

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 > swap** 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 Enterprise Session Border Controller is based on a system-level file called MACTAB that maps interface MAC addresses to interface naming that can be applied within Oracle Communications Enterprise Session Border Controller 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
slp0        00:16:3E:30:00:2D  # Second media interface
=====
```

## interface-mapping

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

Command	Description
interface-mapping show	Display the existing content of /boot/mactab file, with the mapping information of all the available Ethernet Interface Names versus Physical Interface MAC addresses, along with any customer provided label information.
interface-mapping locate <ethernet if name> <seconds>	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.
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>	<p>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.</p> <pre data-bbox="971 1136 1442 1360">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</pre> <p>Then you can execute the interface-mapping swap command.</p> <pre data-bbox="971 1444 1442 1669">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</pre> <p>A warning message appears. Once you confirm the action, the MAC addresses and their corresponding labels are swapped in the /boot/mactab/file.</p>



**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 Enterprise Session Border Controller CLI. 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 Enterprise Session Border Controller 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.



# 7

## Software Upgrade

This section provides information about how to upgrade your Oracle Communications Enterprise Session Border Controller 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 Enterprise Session Border Controller 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 Enterprise Session Border Controller 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 Enterprise Session Border Controller Release Notes for any caveats involving software upgrades.

### Acme Packet 4500 Upgrade Requirements

- To run release Version E-CZ7.5.0 on an Acme Packet 4500 with a transcoding NIU, you must also have a High Speed 4-Fan Module in the chassis.
- To run release Version E-CZ7.5.0 on an Acme Packet 4500, the Stage1 and Stage2 bootloaders must be dated July 3, 2013 or later. See the Oracle Communications Enterprise Session Border Controller Release Notes for more information. You may use the **show version boot** command to confirm.

```
ORACLE# show version boot
Bootloader Info
-----
Stage 1: Jul  3 2013 13:16:30
Stage 2: Jul  3 2013 13:16:30
```

### Stand-alone Upgrade

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

## Check /boot for free space

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

```
ESBC# 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/`.

```
[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]$
```

The Stage3 boot loader is ready for operation after upload and filename change.

## Software Upgrade Procedure

The following procedure describes how to upgrade a Oracle Communications Enterprise Session Border Controller with a new software image. In this procedure, the image file is located on the Oracle Communications Enterprise Session Border Controller'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`.

```
ESBC1# configure terminal  
ESBC1(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 Enterprise Session Border Controller using the **reboot** command.


The Oracle Communications Enterprise Session Border Controller 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.



## HA Upgrade Procedure

1. Confirm that ESBC 1 and ESBC2 start up and are synchronized.

Ensure that the running and current configurations on ESBC1 and ESBC2 have the same number. In the following examples, all of the configuration versions are 5.

On ESBC1 and ESBC2, use the **show health** command to ensure all processes are synchronized.

On ESBC1, show the current configuration version by using the ACLI **display-current-cfg-version** command. Then use the same command on ESBC2 and be sure that its current configuration version is the same as the one on ESBC1.

```
ESBC1# display-current-cfg-version
Current configuration version is 5
ESBC1#
ESBC2# display-current-cfg-version
Current configuration version is 5
ESBC2#
```

On ESBC1, show the running configuration version by using the ACLI **display-running-cfg-version** command. Then use the same command on ESBC2 and be sure that its running configuration version is the same as the one on ESBC1.

```
ESBC1# display-running-cfg-version
Running configuration version is 5
ESBC1#
ESBC2# display-running-cfg-version
Running configuration version is 5
ESBC2#
```

2. On ESBC1, before loading the software image to the flash, check the remaining space in the `/boot` directory using the ACLI **show space boot** command.

```
ESBC1# show space boot
boot: 24759488/25760512 bytes (99%) remaining
ESBC1#
```

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 ESBC 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 ESBC1 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 ESBC1:

- In the ACLI configure terminal menu, type **bootparam** and press <Enter> to bring 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.

```
ESBC1# configure terminal
ESBC1(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 ESBC1.

5. After ESBC1 has completed the boot process, use the **verify-config** command to confirm that the configuration has been upgraded properly.

```
ESBC1# verify-config
```

6. Confirm the ESBC1 is running the new boot image using the **show version** command.

```
ESBC1# show version
Acme Packet 4500 SCZ7.3.0
Build Date=09/09/15
```

7. Use the **show health** command to confirm that ESBC1 is the standby system.
8. As you did for ESBC1, configure the boot parameters on ESBC2 to boot from the new software image. Then reboot ESBC2.

```
ESBC2# reboot
```

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

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

Rebooting ESBC2 causes ESBC1 to become the active system in the HA node.

- When ESBC2 is finished rebooting, use the **show health** command to confirm that it is in the standby state.

 **Note:**

If you need to revert to 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 Enterprise Session Border Controllers 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.

In the descriptions and processes outlined below, ESBC1 is the active system and ESBCOracle Communications Enterprise Session Border Controller2 is the standby system. The procedure uses these designations because when you have completed upgrade process specific to these releases, ESBC1 is the active system.

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

- Change the boot parameters on ESBC2 to use the appropriate Release S-CZ7.2.0 software image.

Using one of these methods, access the boot parameters on ESBC2:

- Reboot the ESBC 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`.

```
ESBC2# configure terminal
ESBC1(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.

```
ESBC(configure)# exit
```

- Reboot ESBC2.

- Using the ACLI **show version** command to confirm that you are using the appropriate release.

```
ESBC2# show version
Acme Packet 4500 SCZ7.2.0
Build Date=01/09/15
```

- Initiate a switchover on ESBC2.

```
ESBC2# notify berpd force
```

At this point, ESBC2 becomes the active system running Release S-CZ7.2.0. ESBC1 is now the standby system running Release S-CZ7.3.0.

- On ESBC1, change the boot parameters as you did in Step 1 of this procedure.
- Reboot ESBC1.

## HA Backout Procedure

If you reach the point in your upgrade procedure where you have upgraded both Oracle Communications Enterprise Session Border Controllers in the HA node 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.

In the descriptions and processes outlined below, Oracle Communications Enterprise Session Border Controller1 is the active system and Oracle Communications Enterprise Session Border Controller2 is the standby system. The procedure uses these designations because when you have completed upgrade process specific to these releases, Oracle Communications Enterprise Session Border Controller1 is the active system.

To backout to a previous (older) release with the both Oracle Communications Enterprise Session Border Controllers in the HA node upgraded:

- Change the boot parameters on Oracle Communications Enterprise Session Border Controller2 to use the appropriate Release SCX6.4.0 software image.

Using one of these methods, access the boot parameters on Oracle Communications Enterprise Session Border Controller2:

- Reboot the Oracle Communications Enterprise Session Border Controller 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 `/code/images/nnSCX640.xz` and `/boot/nnSCZ712.64.bz`.

```
ORACLE2# configure terminal
ORACLE1(configure)# bootparam
'.' = clear field; '-' = go to previous field; ^D = quit
boot device          : eth0
processor number     : 0
host name            : boothost
file name            : /boot/nnSCZ712.64.bz /code/images/nnSCX640.xz
```

In the example above, type the appropriate Release S-CX6.4.0 file name next to the Release S-CZ7.1.2 file name. Press <Enter> to continue scrolling through the boot parameters.

Exit to the main Superuser prompt.

```
ORACLE2(configure)# exit
ORACLE2#
```

2. Reboot Oracle Communications Enterprise Session Border Controller2.
3. Using the ACLI **show version** command to confirm that you are using the appropriate release.

```
ORACLE2# show version
ACME PACKET 4500 Firmware S-CX6.4.0 GA
07/15/10
ORACLE2#
```

4. Initiate a switchover on Oracle Communications Enterprise Session Border Controller2.

```
ORACLE2# notify berpd force
```

At this point, Oracle Communications Enterprise Session Border Controller2 becomes the active system running Release S-CX6.4.0. Oracle Communications Enterprise Session Border Controller1 is now the standby system running Release S-CZ7.1.2

5. On Oracle Communications Enterprise Session Border Controller1, change the boot parameters as you did in Step 1 of this procedure.
6. Reboot Oracle Communications Enterprise Session Border Controller1.

## Moving a Configuration

This section outlines a process for moving an existing Oracle Communications Enterprise Session Border Controller configuration to a new system. Process summary:

1. Create a backup configuration file on the source Oracle Communications Enterprise Session Border Controller.
2. Using SFTP, copy the source backup from the source to the destination Oracle Communications Enterprise Session Border Controller .
3. Restore the newly-transferred backup on the target Oracle Communications Enterprise Session Border Controller.

## Backup Commands

The Oracle Communications Enterprise Session Border Controller software includes a set of commands for easily working with backup configurations. These commands are **backup-config**, **display-backups**, **delete-backup-config**, **restore-backup-config**.

To back up a configuration, use the **backup-config** command. You can confirm that your backup has been created with the **display-backups** command. When the **backup-config** command is executed, the system checks for sufficient resources to complete the operation. If resources are sufficient, the system creates the backup. If resources are insufficient, the task is not completed and the system displays the limiting resources and recommends completing the task at another time.

Backups are created as gzipped tar files in a .tar.gz format. They are stored in the /code/bkups directory on the Acme Packet 4000.



## Backing up the current configuration

To create a backup:

1. In superuser mode, use the **backup-config** command followed by a descriptive filename for the backup you are creating.

```
ORACLE#backup-config 02_Feb_2008
task done
ORACLE#
```

## Listing Backups

You can view the backups available on your system using the **display-backups** command.

To list available backup configurations:

1. In Superuser mode, enter the **display-backups** command. A list of available backup files from the `/code/bkups` directory is displayed on the screen.

```
ORACLE# display-backups
test_config.tar.gz
test-config.tar.gz
runningcfgtest.tar.gz
runningtest_one.tar.gz
BACK_UP_CONFIG.tar.gz
02_Feb_2008.tar.gz
01_Feb_2008.tar.gz
ORACLE#
```

## Copy the Backup to the destination

Send the backup configuration file by way of SFTP from the source to destination Oracle Communications Enterprise Session Border Controller.

To copy a backup configuration from the source to destination Oracle Communications Enterprise Session Border Controller:

1. Use an SFTP client to connect to the ESBC. The management IP address is configured in the `bootparams`.
2. Change directory to where you want to upload a file.
  - `cd /code/bkups` for backup configurations
3. Type `bin` and press Enter to force the SFTP program into binary mode.
4. Upload the file you want to transfer by typing the filename and pressing Enter.

```
C:\Documents and Settings>sftp 172.30.55.127
Connected to 172.30.55.127.
220 VxWorks (1.0) FTP server ready
User (172.30.55.127:(none)): user
331 Password required
Password:
230 User logged in
sftp> cd /code/bkups
250 Changed directory to "/code/bkups"
sftp> bin
200 Type set to I, binary mode
```

```
sftp> put 02_Feb_2008.tar.gz
200 Port set okay
150 Opening BINARY mode data connection
226 Transfer complete
sftp: 9587350 bytes sent in 51.64Seconds 185.65Kbytes/sec.
sftp>
```

## Restoring Backups

To restore a backup configuration on the Oracle Communications Enterprise Session Border Controller:

1. In Superuser mode, enter the **restore-backup-config** command followed by the backup filename you want to restore to the current configuration. You must explicitly name the backup file you want to restore, including the file extension

```
ORACLE# restore-backup-config 02_Feb_2008.tar.gz
Need to perform save-config and activate/reboot activate for changes to take
effect...
task done
ORACLE#
```

2. Correct the Virtual MAC address configuration established on the former device to be suitable for the new device.

Establish the base MAC needed for HA operation by, first, determining the base MAC by way of the ethernet address value of the show media physical command.

```
ORACLE#show media physical
s0p0 (media slot 0, port 0)
  Flags: UP BROADCAST MULTICAST ARP RUNNING
  Type: ETHERNET_CSMACD
  Admin State: enabled
  Auto Negotiation: enabled
...
  Ethernet address is 00:08:25:01:08:44
```

Apply the formula for calculating virtual MAC addressing to the MAC addressing used for this system. This formula is described in the ACLI Configuration Guide.

Configure the physical interfaces with the computed virtual MAC addressing. Refer to the following command line sequence as an example of this procedure.

```
ORACLE# configure terminal
ORACLE(configure)# system
ORACLE(system)# phy-interface
ORACLE(phy-interface)# select
<name>:
1: s0p0
2: s1p0
selection: 1
ORACLE(phy-interface)# virtual-mac 00:08:25:01:08:48
ORACLE(phy-interface)# done
phy-interface
      name                s0p0
  operation-type          Media
      port                0
      slot                0
  virtual-mac              00:08:25:01:08:48
```

3. Save the configuration.

ORACLE# **save-config**

4. Activate the configuration.

ORACLE# **activate-config**

# 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 Enterprise Session Border Controller
- The boot file used to boot the Oracle Communications Enterprise Session Border Controller 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 1 x GbE ports 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 1 x GbE ports 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.



### 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 apegPowerSupply object to include the apegPowerSupply 500 W object.

## Acme Packet 3900 MIBS Paths

Paths for Acme Packet 3900 MIBS.


SNMPv2-SMI::mib-2.47.1.1.1.2.1 = 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 4500 Physical Interfaces

The Acme Packet 4500 supports a single network interface unit (NIU) that contains all external interfaces, including console, alarm, network management and media interfaces. There are multiple NIU types available, which define the supported cabling, speed and operational features, such as transcoding and/or encryption. The NIUs used on the Acme Packet 4500 differs in terms of speed and operational features.

The graphic below shows the location and labeling of the Acme Packet 4500 media and network management ports. This labeling is an important point of reference when you set up the **phy-interface** configuration element.





The Acme Packet 4500 NIU includes the following ports (from left to right).

- Console—Provides serial access for administrative and maintenance purposes.
- Alarm—Dry contact alarm port.
- Mgmt0 to Mgmt2—The system uses these 10/100/1000 Base-T Ethernet ports for device management functions.
- P0,0 to P1,1—The system uses these 4 x GbE 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 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
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


The Acme Packet 4500 NIU supports six kinds of NIUs. Your NIU may support copper, single mode or multimode fiber with an LC connector.

- 4-port GigE copper (RJ45)
- 4-port GigE SFP (LX, SX, or Copper)
- 4-port GigE SFP with QoS and IPSec (LX, SX, or Copper)
- 4-port GigE SFP with IPSec (LX, SX, or Copper)
- 4-port GigE SFP with QoS (LX, SX, or Copper)
- 4-port GigE SFP ETC NIU (LX, SX, or Copper)

## 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.



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
P4	media	0	4
P5	media	0	5

## Acme Packet 6300/6350 Physical Interfaces


The Acme Packet 6300/6350 comes with management interfaces, including console, alarm, USB and network management integrated into the rear of the chassis, above the power supplies. It also supports a two network interface units (NIUs) that contain media interfaces. There is currently one type of media NIU available, which defines the supported cabling and speed.

### Acme Packet 6300/6350 Management Interfaces

Management interfaces include 3 x 10/100/1000 Ethernet interfaces and are labeled Mgmt 0, 1 and 2. (These interfaces are often referred to as wancom0, 1 and 2.) 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.

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

The USB port is reserved for use by Oracle personnel only. There is no front facing console port. The graphic below uses numbers to label the ports.



The table below maps the graphic's reference numbers and the labeling on the management interfaces. It also, when applicable, lists the **operation-type**, **slot** and **port** parameters in the **phy-interface** configuration element for the management interfaces.

Graphic Label	Port 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 6300/6350 Media Interfaces


The Acme Packet 6300/6350 has 3 PHY card slots. The bottom slot, slot0, and the middle slot, slot1, support network-facing media interface cards. The top slot, slot2, supports feature-based hardware, such as transcoding, and is referred to as the "resource" slot. Standard Acme Packet 6300/6350 PHY cards contain 2 x 10-gigabit Ethernet ports.

Ensure that the first 2x10Gig NIU in your system populates slot0. If you are use a second 2x10Gig NIU, insert it into slot1.

#### Note:

Do not insert any 2x10Gig NIU in the resource slot (slot2).

The Acme Packet 6300/6350 supports up to 4 x 10 gigabit media ports each running at full duplex line rate for all packet sizes. The graphic below labels the media ports using numbers.



The table below maps the graphic's reference numbers 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. Note that the media interfaces are not uniquely labeled with the NIU silkscreen. The table distinguishes between these using "left" and "right", with the perspective being the user looking at the NIU panel.

Graphic Label	Port 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

# B

## DPDK-Based OVS Installation

System Design and/or Deployment Engineering may decide to use Open V-Switch (OVS) within Data Path Development Kit (DPDK) architectures. This section explains how to create a DPDK-based OVS Installation. This installation also sets up multi-queues using **dpdkvhostuser** for media traffic. The procedure below is an example; step details vary based on deployment environment.

The procedure below assumes the virtual SBC (vSBC) is installed and ready to start.

1. Install Fedora 23 from a CD or an ISO image.
2. Change the kernel boot parameters by adding the following arguments to the **/etc/default/grub** file:

```
hugepagesz=1G
hugepages=64
default_hugepagesz=1G
isolcpus=1-2
```

3. Update the installed Fedora 23 packages. If your network is behind a proxy, set the value of proxy in **/etc/dnf/dnf.conf**.

```
proxy=http://proxy.example.com
```

4. Update your system.

```
dnf -y update
```

5. Install additional required packages.

```
dnf groupinstall -y "Development Tools"
```

```
dnf install -y hwloc gcc-c++ openssl-devel glib autoconf libtool
glib2-devel zlibrary gtk+-devel gtk3-devel tuncctl kernel-devel
```

6. Set up environment variables.

```
export DPDK_DIR=/usr/src/dpdk
export DPDK_BUILD=$DPDK_DIR/x86_64-native-linuxapp-gcc/
export OVS_DIR=/usr/src/ovs
export DB_SOCKET=/usr/local/var/run/openvswitch/db.sock
export QEMU_DIR=/usr/src/qemu-2.6.0-rc3
```

7. Download and install the following in the **/usr/src** directory.

- dpdk-16.04.tar.gz
- openvswitch-ovs-v1.1.0pre2-10269-gac93328.tar.gz (April 27th snapshot version)
- qemu-2.6.0-rc3.tar.bz2

- a. Change directories to **/usr/src**.

```
cd /usr/src
```

- b. Install DPDK.

```
git clone http://dpdk.org/git/dpdk
cd dpdk
```

```
echo "CONFIG_RTE_BUILD_COMBINE_LIBS=y" >> config/common_linuxapp
make install T=x86_64-native-linuxapp-gcc DESTDIR=install
cd ..
```

 **Note:**

If make fails because it cannot find `/lib/modules/<version>/build`, run `dnf -y install kernel-devel-<version>` and run make again.

```
dnf -y install kernel-devel-4.5.5-300.fc24.x86_64
```

**c. Install OVS.**

```
git clone https://github.com/openvswitch/ovs
cd ovs
./boot.sh
./configure --with-dpdk=$DPDK_BUILD CFLAGS="-g -O2 -Wno-cast-align"
make install
cd ..
```

 **Note:**

If you get the error message `ImportError: No module named six`, make sure your `http_proxy` variable is set and run `pip install six --proxy=$http_proxy`.

 **Note:**

If `configure` returns an error, see the `INSTALL` file for your distribution.

**d. Install qemu.**

```
wget http://wiki.qemu.org/download/qemu-2.6.0-rc3.tar.bz2
tar jxf qemu-2.6.0-rc3.tar.bz2
cd qemu-2.6.0-rc3
./configure
make install
cd ..
```

**8. Install DPDK modules.**

```
modprobe uio insmod $DPDK_BUILD/kmod/igb_uio.ko
```

**9. Mount the hugepages.**

```
mount -t hugetlbfs -o pagesize=1G none /dev/hugepages
```

**10. Choose the host NICs for guest VM's media traffic. These interfaces need to be bound to DPDK. For example, you may choose `eno1` as the vSBC's access interfaces, `s0p0`.**

Do not use host DPDK interfaces as vSBC wancom interfaces.

**11. Assign these media traffic NICs to OVS DPDK.**

```
$DPDK_DIR/tools/dpdk_nic_bind.py --bind=igb_uio eno1
$DPDK_DIR/tools/dpdk_nic_bind.py --bind=igb_uio eno2
```

**12. Create OVS directories.**

```
mkdir -p /usr/local/etc/openvswitch
mkdir -p /usr/local/var/run/openvswitch
rm -f /usr/local/etc/openvswitch/conf.d
```

**13. Create the OVS database schema.**

```
ovsdb-tool create /usr/local/etc/openvswitch/conf.db /usr/local/share/
openvswitch/vswitch.ovsschema
```

**14. Start the ovsdb server/ovsvswitch.**

```
ovsdb-server --remote=punix:$DB_SOCKET --
remote=db:Open_vSwitch,Open_vSwitch,manager_options --private-
key=db:Open_vSwitch,SSL,private_key --
certificate=Open_vSwitch,SSL,certificate --bootstrap-ca-
cert=db:Open_vSwitch,SSL,ca_cert --pidfile --detach

ovs-vsctl --no-wait init

ovs-vswitchd --dpdk -vhost_sock_dir /usr/local/var/run/openvswitch -c 0xe -
n 4 --socket-mem 16384,0 -- unix:$DB_SOCKET --pidfile --detach --log-file=/usr/
local/var/log/openvswitch/ovs-vswitchd.log
```

**15. Set up the CPU mask so that these cores are used for pmd threads in the ovsvswitch.**

```
ovs-vsctl set Open_vSwitch . other_config:pmd-cpu-mask=6
```

This starts two PMD threads, and therefore two RX/TX queue pairs.

**16. Create bridges and ports. (You can create multiple ports to launch multiple guest VMs.)**

```
ovs-vsctl add-br enolbr -- set bridge enolbr datapath_type=netdev
ovs-vsctl add-port enolbr dpdk0 -- set Interface dpdk0 type=dpdk
ovs-vsctl add-port enolbr vhost-user-port0-on-enolbr -- set Interface vhost-
user-port0-on-enolbr type=dpdkvhostuser
ovs-vsctl set Interface vhost-user-port0-on-enolbr options:n_rxq=2

ovs-vsctl add-br eno2br -- set bridge eno2br datapath_type=netdev
ovs-vsctl add-port eno2br dpdk1 -- set Interface dpdk1 type=dpdk
ovs-vsctl add-port eno2br vhost-user-port0-on-eno2br -- set Interface vhost-
user-port0-on-eno2br type=dpdkvhostuser
ovs-vsctl set Interface vhost-user-port0-on-eno2br options:n_rxq=2

ovs-vsctl add-br ens1f2br -- set bridge ens1f2br datapath_type=netdev
ovs-vsctl add-port ens1f2br ens1f2
ip tuntap add mode tap tap0-on-ens1f2br
ip link set dev tap0-on-ens1f2br up
ovs-vsctl add-port ens1f2br tap0-on-ens1f2br
```

**17. Export your display.**

```
export DISPLAY=IPADDRESS_WHERE_YOU_WANT_TO_SEE_QEMU_GRAPHICAL_OUTPUT:1
```

**18. Launch your VM.**

```
/usr/local/bin/qemu-system-x86_64 -name vsbcl -S -machine pc-
i440fx-2.1,accel=kvm,usb=off -cpu SandyBridge -m 8192 -realtime mlock=off \
-smp 4,sockets=1,cores=4 \
-rtc base=utc -no-shutdown -boot strict=on \
-device piix3-usb-uhci,id=usb,bus=pci.0,addr=0x1.0x2 \
-device virtio-serial-pci,id=virtio-serial0,bus=pci.0,addr=0x5 \
-drive file=/usr/src/vms/nnSCZ739a6-img-vm_kvm.qcow2,if=none,id=drive-
ide0-0-0,format=qcow2,cache=none \
```

```

-device ide-hd,bus=ide.0,unit=0,drive=drive-ide0-0-0,id=ide0-0-0,bootindex=1
\
-device virtio-balloon-pci,id=balloon0,bus=pci.0,addr=0x6 \
-monitor stdio \
-netdev tap,script=no,id=mgmt,ifname=tap0-on-ens1f2br \
-device e1000,netdev=mgmt,mac=52:55:00:dd:ee:01 \
-chardev socket,id=char1,path=/usr/local/var/run/openvswitch/vhost-user-
port0-on-eno1br \
-netdev type=vhost-user,id=access,chardev=char1,vhostforce=on,queues=2 \
-device virtio-net-pci,netdev=access,mac=52:55:00:dd:ee:02,mq=on,vectors=6 \
-chardev socket,id=char2,path=/usr/local/var/run/openvswitch/vhost-user-
port0-on-eno2br \
-netdev type=vhost-user,id=core,chardev=char2,vhostforce=on,queues=2 \
-device virtio-net-pci,netdev=core,mac=52:55:00:dd:ee:03,mq=on,vectors=6 \
-object memory-backend-file,id=mem,size=8192M,mem-path=/dev/
hugepages,share=on \
-numa node,memdev=mem -mem-prealloc

```

You see the VM display in a different window. You launch VMs from your qemu control monitor.

19. Execute the **'cont'** command in the qemu control monitor to boot the vSBC in the guest VM.

```
command 'cont'
```