Oracle® Communications Session Border Controller Platform Preparation and Installation Guide





Oracle Communications Session Border Controller Platform Preparation and Installation Guide, Release S-Cz10.0.0 - for Service Provider and Enterprise

G20486-02

Copyright © 2025, 2025, Oracle and/or its affiliates.

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

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

If this is software, software documentation, data (as defined in the Federal Acquisition Regulation), or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then the following notice is applicable:

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

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

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

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

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

Contents

About This Guide	
My Oracle Support	Vİ
Revision History	
Introduction to Platform Preparation and Software De	eployment
Acme Packet Engineered Platforms	1-:
Virtual Machines	1-
Oracle Servers	1-
Multi-NUMA Node Environments	1-7
Oracle Server X8-2 Platform Preparation	
Available Connections	2-
Cable the Oracle X8-2	2-
Cable the Local Console	2-1
Connect ILOM to the Oracle X8-2	2-1
Install the Software on the X8-2	2-12
Next Steps After the Software Installation	2-1
Oracle Server X9-2 Platform Preparation	
Available Connections	3-:
X9-2 Back Panel Connectors and Ports	3-
Cable the Oracle Server X9-2	3-
Cable the Local Console	3-
Connect ILOM to the Oracle Server X9-2	3-
Install the Software on the Oracle Server X9-2	3-
Next Steps After the Software Installation	3-



Session Delivery Products as Virtual Machines 4 Hyperthreading and CPU Affinity 4-1 Host Hypervisor CPU Affinity (Pinning) 4-2 4-3 Configuration Overview **VLAN Support** 4-4 **Provisioning Entitlements** 4-4 5 Virtual Machine Platforms Create and Deploy on KVM 5-1 Create and Deploy on VMware® 5-3 Optimize for VMware 5-4 Create and Deploy on Hyper-V® 5-6 OpenStack Heat Template 5-10 **Public Cloud Platforms** 6 Deploying the SBC on Cloud Infrastructures in Standalone Mode 6-1 Deploying the SBC on Cloud Infrastructures in HA Mode 6-2 Key SBC Configuration and Operation Detail in HA Mode 6-3 Cloud-Based HA Report Data 6-6 6-7 Cloud-Specific HA Deployment Considerations HA over OCI Overview 6-7 6-7 HA over AWS Overview Create and Deploy SBC Instances on EC2 6-8 6-8 Prerequisites to EC2 Deployment Generate an EC2 AMI from the vSBC Image 6-9 EC2 Deployment Procedure 6-10 Deploy AWS AMI with Terraform 6-19 6-25 Create and Deploy on OCI Prerequisites to Deploying an OCI Instance 6-25 Create Dynamic Group and Policy Statements 6-27 Deploying the OCI Instance 6-28 Migrating to an Optimized3.Flex Machine Shape 6-34 Create and Deploy on OCI using Resource Manager 6-36 Prerequisites 6-37 6-38 Creating VCN Stack 6-43 Creating SBC Stack **Destroying Resources** 6-51 Running templates from CLI 6-51



IAM Policies

6-53

	Troubleshooting	6-54
	Create and Deploy on Azure	6-55
	Prerequisites to Deploying an Azure Instance	6-55
	Deploying the Azure Instance	6-57
	Create Networking for Additional Interfaces	6-61
	Set the User and Administrative Passwords on the OCSBC	6-62
	Deploying on Azure with accelerated networking	6-63
	Create and Deploy on Google Cloud Platform	6-64
	Prerequisites to Deploying an GCP Instance	6-65
	Deploy a Standalone SBC	6-69
	Deploy High Availability SBCs	6-69
	Configure an HA Pair on GCP	6-72
7	Boot Management	
	Boot Loader Overview	7-1
	Boot Parameters	7-2
	Boot Parameter Definitions	7-3
	Boot Flags	7-4
	Changing Boot Parameters	7-4
	Change Boot Parameters from the ACLI	7-5
	Change Boot Parameters by Interrupting a Boot in Progress	7-6
	Displaying Files from Boot and Images directory	7-7
8	Formatting the Disk Volume	
	40GB or Less Format Plan	8-1
	40GB or More Default Format Plan	8-1
	40GB or More Custom Format Plan	8-1
	Formatting Procedure	8-2
9	Interface Considerations for VM Platforms	
	Software Adaptations for Physical Interfaces	9-1
	COTS Interfaces	9-1
	COTS Network Interfaces	9-1
	The interface-mapping Branch	9-2
	Working with the interface-mapping branch	9-2
	interface-mapping	9-3
	Serial Interfaces	9-4



10 Flash Drive Installation via Boot Media Creator Creating a Build Image 10-1 11 Software Upgrade **Upgrade Checklist** 11-1 Check Upgrade Readiness 11-2 Download the Software 11-4 Check /boot for free space 11-4 Select Upgrade Procedure 11-5 Upload the Stage 3 Boot Loader and System Image 11-5 Stand-alone Upgrade 11-6 **HA** Upgrade 11-7 HA Upgrade Procedure 11-8 **HA Backout Procedure** 11-10 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 3950/4900 Physical Interfaces A-5 Acme Packet 4600 Physical Interfaces A-6 Acme Packet 6100 Physical Interfaces A-7 Acme Packet 6300/6350 Physical Interfaces A-8 Signaling and Media Interfaces A-9 Acme Packet 6300 A-10 Acme Packet 6350 A-11 Dual Port and Quad Port NIUs A-11 Acme Packet 6400 Physical Interfaces A-14 VNF Metadata and Userdata Example



B

About This Guide

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

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

This publication is used with Oracle Communications Session Border Controller and Oracle Enterprise Session Border Controller.

Documentation Set

The following table describes the documentation set for this release.

Document Name	Document Description
Acme Packet 3900 Hardware Installation Guide	Contains information about the components and installation of the Acme Packet 3900.
Acme Packet 4600 Hardware Installation Guide	Contains information about the components and installation of the Acme Packet 4600.
Acme Packet 4900 Hardware Installation Guide	Contains information about the components and installation of the Acme Packet 3950 and Acme Packet 4900.
Acme Packet 6350 Hardware Installation Guide	Contains information about the components and installation of the Acme Packet 6350.
Acme Packet 6400 Hardware Installation Guide	Contains information about the components and installation of the Acme Packet 6400.
Release Notes	Contains information about the current documentation set release, including new features and management changes.
Known Issues & Caveats	Contains known issues and caveats
Configuration Guide	Contains information about the administration and software configuration of the Service Provider Session Border Controller (SBC).
ACLI Reference Guide	Contains explanations of how to use the ACLI, as an alphabetical listings and descriptions of all ACLI commands and configuration parameters.
Maintenance and Troubleshooting Guide	Contains information about SBC logs, performance announcements, system management, inventory management, upgrades, working with configurations, and managing backups and archives.



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

Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc.

My Oracle Support

My Oracle Support (https://support.oracle.com) is your initial point of contact for all product support and training needs. A representative at Customer Access Support (CAS) can assist you with My Oracle Support registration.

Call the CAS main number at 1-800-223-1711 (toll-free in the US), or call the Oracle Support hotline for your local country from the list at http://www.oracle.com/us/support/contact/ index.html. When calling, make the selections in the sequence shown below on the Support telephone menu:

- Select 2 for New Service Request.
- 2. Select 3 for Hardware, Networking, and Solaris Operating System Support.
- Select one of the following options:
 - For technical issues such as creating a new Service Request (SR), select 1.
 - For non-technical issues such as registration or assistance with My Oracle Support, select 2.



You are connected to a live agent who can assist you with My Oracle Support registration and opening a support ticket.

My Oracle Support is available 24 hours a day, 7 days a week, 365 days a year.

Emergency Response

In the event of a critical service situation, emergency response is offered by the Customer Access Support (CAS) main number at 1-800-223-1711 (toll-free in the US), or call the Oracle Support hotline for your local country from the list at http://www.oracle.com/us/support/contact/index.html. The emergency response provides immediate coverage, automatic escalation, and other features to ensure that the critical situation is resolved as rapidly as possible.

A critical situation is defined as a problem with the installed equipment that severely affects service, traffic, or maintenance capabilities, and requires immediate corrective action. Critical situations affect service and/or system operation resulting in one or several of these situations:

- A total system failure that results in loss of all transaction processing capability
- Significant reduction in system capacity or traffic handling capability
- Loss of the system's ability to perform automatic system reconfiguration
- Inability to restart a processor or the system
- Corruption of system databases that requires service affecting corrective actions
- Loss of access for maintenance or recovery operations
- Loss of the system ability to provide any required critical or major trouble notification

Any other problem severely affecting service, capacity/traffic, billing, and maintenance capabilities may be defined as critical by prior discussion and agreement with Oracle.

Locate Product Documentation on the Oracle Help Center Site

Oracle Communications customer documentation is available on the web at the Oracle Help Center (OHC) site, http://docs.oracle.com. You do not have to register to access these documents. Viewing these files requires Adobe Acrobat Reader, which can be downloaded at http://www.adobe.com.

- 1. Access the Oracle Help Center site at http://docs.oracle.com.
- Click Industries.
- 3. Under the Oracle Communications sub-header, click the **Oracle Communications** documentation link.

The Communications Documentation page appears. Most products covered by these documentation sets appear under the headings "Network Session Delivery and Control Infrastructure" or "Platforms."

- Click on your Product and then Release Number.
 A list of the entire documentation set for the selected product and release appears.
- 5. To download a file to your location, right-click the **PDF** link, select **Save target as** (or similar command based on your browser), and save to a local folder.

Access to Oracle Support

Oracle customers that have purchased support have access to electronic support through My Oracle Support. For information, visit http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info or visit http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs if you are hearing impaired.



Revision History

The following table shows the dates and descriptions of revisions to the Platform Preparation and Installation Guide.

Date	Description	
March 2025	Initial release.	
April 2025	 Adds AP6400 section to appendix 	



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.
- 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.

Oracle Servers

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

Multi-NUMA Node Environments

In a multi-NUMA node environment, SBC performance can degrade significantly if traffic processing tasks are not properly bound to the same NUMA node. The penalties associated with inter-NUMA communication primarily arise due to increased memory access latency and cross-socket CPU communication overhead.

Sources of Inter-NUMA Node Penalties include:

Remote Memory Access Latency
 When an SBC process running on NUMA Node 0 accesses memory allocated on NUMA
 Node 1, it introduces higher latency compared to local memory access.

The latency can be 30-50% higher or more when accessing remote memory compared to local memory.

Cross-NUMA CPU Synchronization
 If different SBC threads (e.g., signaling and media handling) are spread across NUMA nodes, the L3 cache coherence traffic increases across sockets.

This leads to higher interconnect bandwidth usage (e.g., Intel UPI or AMD Infinity Fabric), causing potential performance bottlenecks.

 Network Interrupt Processing Across NUMA Nodes
 If a NIC is assigned to NUMA Node 0, but the SBC process handling packets is running on NUMA Node 1, then every packet transfer incurs an inter-socket latency.

This results in degraded packet processing performance, especially in high-throughput VoIP and SIP traffic environments.



Oracle Server X8-2 Platform Preparation

Oracle Communications produces a variety of software products that run on the Oracle Server X8-2 platform. See the Release Notes for which Oracle Communications applications run on the X8-2.

Use your Hardware documentation to install and establish system management by way of Oracle Integrated Lights Out Manager (ILOM). Then use the steps below to prepare the Oracle X8-2 for session delivery software installation.

Note:

The ILOM Cable Connection procedure also displays ILOM cabling.

- Confirm applicable firmware on the server.
 - To check the firmware versions installed in the server, go to the ILOM web interface, and navigate to System Information, Firmware.
 - Software and firmware versions qualified for use with Oracle Session Delivery products include:
 - ILOM—v4.0.3.34
 - BIOS—51.01.01.00
- 2. Upgrade or downgrade the firmware on the server as necessary. Go to https://docs.oracle.com/cd/E81115 01/index.html for ILOM upgrade instructions.
- 3. Configure the BIOS settings. (Settings navigation may differ based on the BIOS version.)
 - a. Observe the boot procedure, logged to the console during bootup, and use the documented key sequence to interrupt the boot and display the BIOS configuration dialogs. For example, pressing the F2 key is a common way to enter BIOS configuration from a terminal application that supports function keys.
 - **b.** Navigate to the Boot menu and, depending on the software distribution you are using, set the USB or CD as the first device followed by the disk controller. (Navigation: Boot)
 - **c.** Disable Hyper-Threading. (Navigation: Advanced, Processor Configuration, Hyper-Threading)

Note:

Refer to Hyperthreading and CPU Affinity for Oracle guidelines on the use of Hyper-threading.

- d. Disable CPU power limit. (Navigation: Advanced / CPU Power Management Configuration)
- e. Disable C6 Reporting. (Navigation: Advanced / CPU Power Management Configuration, CPU C6 report)


- f. Change Energy Performance to Performance. For example, set "ENERY_PERF_BIAS_CFG" mode to "PERF". (Navigation: Advanced / CPU Power Management Configuration, Energy Performance)
- g. To decrease boot up time, Oracle recommends disabling Intel PXE Boot Agent for both onboard and NIC ethernet ports. Press F2 and navigate to Advanced, Network Stack Configuration. Then disable IPv4 PXE support.



PXE boot is not supported in this release.

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


Figure 2-1 Selecting RAID Configuration



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




Figure 2-2 Begin RAID Configuration



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

Figure 2-3 Clear Any Existing RAID Configuration






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

Figure 2-4 RAID - Create Virtual Drive




e. Set the RAID level to RAID-1.

Figure 2-5 Set Drive to RAID1




f. Select your drives.

Figure 2-6 RAID - Select Drives




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

Figure 2-7 Select All Drives



h. Save the RAID configuration.


Figure 2-8 Save RAID Configuration



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




Figure 2-9 Initialize RAID Configuration



j. After the initialization completes, return to the Main Menu to Save and Exit.

Figure 2-10 Exit RAID Configuration



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

Available Connections

Please read all of the information for each of the available connections prior to cabling the Oracle X8-2.

Port	Description	You Need:
NET (0-2)	From left to right: 1 GigE ports - Net 0 10 GigE ports - Net 1, Net 2 Enables you to connect the X8-2 to your network.	A Category 6 (or better) Ethernet cable to connect to the NET 0 port to your network Network parameters such as an IP address (can be provided by DHCP services or assigned a static address in the OS) Additional Category 6 (or better) Ethernet cables and Ethernet addresses as needed for additional connections to NET 0, 1 and 2.
NET MGT	Provides a 10/100/1000 BASE-T Ethernet connection to the Service Processor (SP) through an RJ-45 connector. The NET MGT port provides support connections to the SP using the Oracle Integrated Lights Out Manager (ILOM) CLI and Web interface. By default, the NET MGT port is configured to use DHCP to automatically obtain an IP address. Alternatively, you can assign a static IP address to the NET MGT port, you must configure its network settings. When configured, use the NET MGT port IP address to log on to the device using a browser or secure shell.	Category 6 (or better) Ethernet cable to connect the NET MGT port to your network IP address for this port (required from DHCP or a static address)
SER MGT (COM1)	Provides a TIA/EIA-232 serial	A terminal device (For example, terminal, connection to a terminal server, or computer such as a laptop running terminal emulation software) A cable to connect the terminal device to the SER MGT (COM1) port



Port	Description	You Need:
USB	Provides USB3.0 connection to the computer. You can connect and disconnect USB cables to the USB port without affecting server operations.	Installation media Note: Maximum USB cable length: 5 meters

Cable the Oracle X8-2

After mounting the Oracle X8-2 in an equipment rack and installing all components, use the following instructions to connect all appropriate data cables to the ports before powering the system up and beginning the configuration.

Oracle qualified the following configurations of the Oracle X8-2.

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



The X8-2 does not support the 40G interface speed.

On board interfaces for all configurations include:

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



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



Figure 2-11 Oracle X8-2 Configuration A (4x10 GigE NIC)




Figure 2-12 Oracle X8-2 Configuration B (Two 4x10 GigE NICs)





Figure 2-13 Oracle X8-2 Configuration B (One QSFP and One 4x10 GigE NICs)



Caution:

Please review your Oracle X8-2 Product Notes. Notes for release 1.1.2 describes physical issues with some optical transceivers installed into an SFP28 port.

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

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

Cable the Local Console

You can connect the Administration console to the local SER MGT (COM1) serial console port. You can also operate the console using serial emulation over ILOM.

To cable a console connection:

- Connect a serial console cable with an RJ-45 connector to the SER MGT port.
- Connect ethernet to the ILOM port and use serial emulation.

When configuring boot loader parameters, set the **Console Device** to COM1. Never use COM2 or VGA. The Oracle X8-2 server cannot boot the system when set to the default of VGA. You must change this bootparameter when deploying over this platform.

Refer to the section "Change Boot Parameters by Interrupting a Boot in Progress" within the *Installation and Platform Preparation Guide* to learn how to set your **Console Device** bootparameter to "COM1". Refer to the section Set Up a Remote Console (SSH) to learn how to run an SSH session via iLOM using Virtual Serial Port Emulation.

Follow this procedure to cable your console:

- 1. Locate the appropriate cables to connect to the Oracle X8-2.
- To cable a serial connection, insert the serial console cable into the SER MGT (COM1) port.

Figure 2-14 Connecting to USB and SER MGT (COM1) Ports





Refer to the Oracle X8-2 hardware documentation for information on how to configure the terminal application to connect to the console, and how to establish communications with the Oracle X8-2.

- 3. For installation procedures, insert the USB stick in the USB port.
- 4. Lead the cables neatly away from the rear panel.

5. Plug in the cables to their respective destination components.

Connect ILOM to the Oracle X8-2

Use the following procedure to make a connection to the Oracle X8-2 Oracle Integrated Lights Out Manager (ILOM) port. For a remote permanent connection to the Service Processor over the ILOM connection, use the rear panel NET MGT port.

Note:

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

- Category 6 (or better) Ethernet
- Locate the cable to connect to the Oracle X8-2 for Communications.
- 2. Plug the RJ-45 connector into the ILOM port.

Figure 2-15 Connecting to ILOM over the Network



- 3. Lead the cable neatly away from the rear panel.
- 4. Connect the other end of the cable to the LAN.
- Refer to the Oracle X8-2 hardware documentation for information about how to configure the Web browser application to connect to the console, and how to establish communications with the Oracle X8-2.

Install the Software on the X8-2

The Oracle Communications Session Router requires software installation when deployed on the Oracle X8-2.

Software installation to Oracle X8-2 includes the following high-level steps:

- Insert your installation media into the USB slot or connect the ISO image by way of Oracle Integrated Lights Out Manager (ILOM) virtual media.
- 2. Power on the Oracle X8-2.
- 3. Observe the startup process, and press F8 to enter the boot menu when it becomes available.
- Select the bootable USB or ISO setting.



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

Save and exit the boot menu.

The Oracle X8-2 starts the Session Router installation.

- 6. Change the Console Device boot parameter to COM1 during installation. If you miss this change during the installation, power on and off the device or catch the boot parameter interrupt and change as soon as possible.
- 7. Remove the USB media when prompted by the Oracle X8-2.
- Allow the Oracle X8-2 complete the installation process and boot to the newly installed Session Router software.

Next Steps After the Software Installation

Oracle recommends the following steps after installation on the Oracle X8-2 platform.

- 1. Execute the **format hard-disk** command, per your requirements. See the "Formatting the Disk Volume" for reference and instructions.
- 2. Turn off the SBC using the **Halt** command. This provides a graceful software shutdown, after which the hardware is still powered on.
- Power cycle the hardware using the power switch, a power controller, or by physically disconnecting and reconnecting the power cable.

To configure the SBC, refer to the ACLI Configuration Guide.

Boot parameter changes to consider prior to service configuration include:

- Set the **Target Name** to your preferred name.
- Set the Console Device to COM1 (serial).
- Set the **IP Address** to your preferred management port IP address.
- Set the Netmask for your management port IP address.
- Set the Gateway address for your management port IP address.

Note:

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



Oracle Server X9-2 Platform Preparation

Oracle Communications produces a variety of software products that run on the Oracle Server X9-2 platform. See the Release Notes for which Oracle Communications applications run on the X9-2.

Use your Hardware documentation to install and establish system management by way of Oracle Integrated Lights Out Manager (ILOM). Then use the steps below to prepare the Oracle X9-2 for session delivery software installation.

- Configure the BIOS settings. See the X9-2 Service Manual for instructions on how to configure each setting.
 - a. Disable Hyper-Threading.



Refer to Hyperthreading and CPU Affinity for Oracle guidelines on the use of Hyper-threading.

- b. Disable C6 Reporting.
- c. Disabling Intel PXE Boot Agent.
- d. Reboot the server.
- 2. Perform a cold shutdown by removing all system power.

Available Connections


Please read all of the information for each of the available connections prior to cabling the Oracle Server X9-2.

Port	Description	You Need:
NET MGT	Provides a 10/100/1000 BASE-T Ethernet connection to the Service Processor (SP) through an RJ-45 connector. The NET MGT port provides support connections to the SP using the Oracle Integrated Lights Out Manager (ILOM) CLI and Web interface. By default, the NET MGT port is configured to use DHCP to automatically obtain an IP address. Alternatively, you can assign a static IP address to the NET MGT port. To use the NET MGT port, you must configure its network settings. When configured, use the NET MGT port IP address to log on to the device using a browser or secure shell.	Category 6 (or better) Ethernet cable to connect the NET MGT port to your network IP address for this port (required from DHCP or a static address)

Port	Description	You Need:
NET 0	The 1 Gbps host management RJ-45 connector port enables you to connect the Oracle Server X9-2 to your network.	A Category 6 (or better) Ethernet cable to connect to the NET 0 port to your network Network parameters such as an IP address (can be provided by DHCP services or assigned a static address in the OS)
SER MGT (COM1)	Provides a TIA/EIA-232 serial Oracle/ Cisco standard connection to the SP through an RJ-45 connector. SER MGT (COM1) connects to either Service Processor by default, but can be redirected to the host.	A terminal device (For example, terminal, connection to a terminal server, or computer such as a laptop running terminal emulation software) A cable to connect the terminal device to the SER MGT (COM1) port
	Default settings: 8N1: eight data bits, no parity, one stop bit 115200 baud	
USB	Provides USB3.0 connection to the computer. You can connect and disconnect USB cables to the USB port without affecting server operations.	Installation media Note: Maximum USB cable length: 5 meters

X9-2 Back Panel Connectors and Ports

The following figure shows the locations of cable connectors and ports on the back of Oracle Server X9-2 and the cables and devices that you connect to them.



Call Out	Cable Port or Expansion Slot	Description	
1	Power supply 0 input power Power supply 1 input power	The server has two power supply connectors, one for each power supply, labeled PS0 and PS1. Power supply 0 input power and Power supply 1 input power both connect to a rack power distribution unit (PDU).	
		Do not attach power cables to the power supplies until you finish connecting the data cables to the server. The server goes into Standby power mode, and the Oracle ILOM service processor initializes when the AC power cables are connected to the power source. System messages might be lost after 60 seconds if the server is not connected to a terminal, PC, or workstation.	
		Oracle ILOM signals a fault on any installed power supply that is not connected to an AC power source, which might indicate a loss of redundancy.	
2	OCP-V3 NIC QSFP	(Optional) 10/25/50/100/200 Gbs Open Compute Project (OCP) Version 3.0 (V3) Network Interface Card (NIC) with two QSFP ports (PORT 1 and PORT 2)	
		Two QSFP 28/56 GbE Ethernet connectors for the Ethernet controller.	
		Note:	
		These ports are not used for the Session Router or Enterprise Session Router.	
3	Network management port (NET MGT)	The service processor NET MGT port is the optional connection to the Oracle ILOM service processor. The service processor NET MGT port uses an RJ-45 cable for a 100/1000BASE-T connection.	
		The NET MGT port is configured by default to use Dynamic Host Configuration Protocol (DHCP).	
4	Host Management Ethernet port	NET 0: 1 Gbps Host Management RJ-45 connector port	
	(NET 0)	The host management Ethernet port enables you to connect the system to the network. The Ethernet port uses an RJ-45 cable for a 1 Gbps Host Management connection.	
5	Serial management port (SER MGT)	The service processor SER MGT port uses an RJ-45 cable and terminal (or emulator) to provide access to the Oracle ILOM command-line interface (CLI). Using Oracle ILOM, you can configure it to connect to the system console.	
		This port does not support network connections.	
6	USB port	The USB port supports hot-plugging. You can connect and disconnect a USB cable or a peripheral device while the server is running without affecting system operations.	

Cable the Oracle Server X9-2

After mounting the Oracle Server X9-2 in an equipment rack and installing all components, use the following instructions to connect all appropriate data cables to the ports before powering the system up and beginning the configuration.

Oracle qualified the following configurations of the Oracle Server X9-2.

- Configuration A: A four-port 10GBASE-T Ethernet NIC
- Configuration B: A four-port 10-Gigabit QSFP+ NIC

The quad-port NICs use Intel XL710 series cards.

On board interfaces for all configurations include:

- One 100/1000 BASE-T RJ-45 Oracle Integrated Lights Out Manager (ILOM) service processor (SP) network management (NET MGT) port
- One 1 Gbps Host Management RJ-45 connector port, labeled NET 0
- One RJ-45 serial management (SER MGT) port

Figure 3-1 Oracle Server X9-2 Configuration A: four-port 10GBASE-T Ethernet NIC




Figure 3-2 Oracle Server X9-2 Configuration B: four-port 10-Gigabit QSFP+ NIC



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

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

After booting the Session Router, use the interface-mapping command to map physical interfaces to configuration names.



Cable the Local Console

You can connect the Administration console to the local SER MGT (COM1) serial console port. You can also operate the console using serial emulation over ILOM.

To cable a console connection:

- Connect a serial console cable with an RJ-45 connector to the SER MGT port.
- Connect ethernet to the ILOM port and use serial emulation.

When configuring boot loader parameters, set the **Console Device** to COM1. Never use COM2 or VGA. The Oracle Server X9-2 server cannot boot the system when set to the default of VGA. You must change this bootparameter when deploying over this platform.

Refer to the section "Change Boot Parameters by Interrupting a Boot in Progress" within the *Installation and Platform Preparation Guide* to learn how to set your **Console Device** bootparameter to "COM1". Refer to Set Up a Remote Console to learn how to run an SSH session via iLOM using Virtual Serial Port Emulation.

Follow this procedure to cable your console:

- 1. Locate the appropriate cables to connect to the Oracle Server X9-2.
- To cable a serial connection, insert the serial console cable into the SER MGT (COM1) port.

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





Refer to the Oracle Server X9-2 hardware documentation for information on how to configure the terminal application to connect to the console, and how to establish communications with the Oracle Server X9-2.

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



Connect ILOM to the Oracle Server X9-2

Use the following procedure to make a connection to the Oracle Server X9-2 Oracle Integrated Lights Out Manager (ILOM) port. For a remote permanent connection to the Service Processor over the ILOM connection, use the rear panel NET MGT port.

Note:

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

- Category 6 (or better) Ethernet
- Locate the cable to connect to the Oracle Server X9-2 for Communications.
- 2. Plug the RJ-45 connector into the ILOM port.

Figure 3-4 Connecting to ILOM over the Network



- 3. Lead the cable neatly away from the rear panel.
- 4. Connect the other end of the cable to the LAN.
- Refer to the Oracle Server X9-2 hardware documentation for information about how to configure the Web browser application to connect to the console, and how to establish communications with the Oracle Server X9-2.

Install the Software on the Oracle Server X9-2

The Oracle Communications Session Router requires software installation when deployed on the Oracle Server X9-2.

Software installation to Oracle Server X9-2 includes the following high-level steps:

- Insert your installation media into the USB slot or connect the ISO image by way of Oracle Integrated Lights Out Manager (ILOM) virtual media.
- 2. Power on the Oracle Server X9-2.
- 3. Observe the startup process, and press F2 to enter the boot menu when it becomes available.
- Select the bootable USB or ISO setting.



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

Save and exit the boot menu.

The Oracle Server X9-2 starts the Session Router installation.

- 6. Change the Console Device boot parameter to COM1 during installation. If you miss this change during the installation, power on and off the device or catch the boot parameter interrupt and change as soon as possible.
- 7. Remove the USB media when prompted by the Oracle Server X9-2.
- 8. Allow the Oracle Server X9-2 complete the installation process and boot /dev/sda to the newly installed Session Router software.

Next Steps After the Software Installation

Oracle recommends the following steps after installation on the Oracle Server X9-2 platform.

- **1.** Execute the **format hard-disk** command, per your requirements. See the "Formatting the Disk Volume" for reference and instructions. .
- 2. Turn off the SBC using the **Halt** command. This provides a graceful software shutdown, after which the hardware is still powered on.
- Power cycle the hardware using the power switch, a power controller, or by physically disconnecting and reconnecting the power cable.

To configure the SBC, refer to the ACLI Configuration Guide.

Boot parameter changes to consider prior to service configuration include:

- Set the **Target Name** to your preferred name.
- Set the Console Device to COM1 (serial).
- Set the IP Address to your preferred management port IP address.
- Set the Netmask for your management port IP address.
- Set the Gateway address for your management port IP address.

Note:

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



4

Session Delivery Products as Virtual Machines

You can deploy version S-Cz10.0.0 of Oracle's Session Delivery Products as Virtual Machine (VM). This document refers to Session Delivery Products generically. See your software version's Release Notes to verify your product's support for deployment as a virtual machine.

Support for the SBC and Session Router as VMs is the same. This document refers to the SBC, but applies equally to the Session Router.

VM deployment types include:

- A standalone (not orchestrated) instance Oracle Communications Session Border Controller operating as a virtual machine running on a hypervisor
- Virtual Machine(s) deployed within an Orchestrated Network Functions Virtualization (NFV) environment
- Virtual Machine(s) deployed within private or public Cloud environments

Standalone SBC VM deployment instances supported for all platforms. Support within an orchestrated environment is dependent on orchestrator and SBC version. High Availability configurations are supported by both deployment types.

You can enable the SBC to utilize CPU hyperthreading, also called SMT. Considerations include your environment's support of the feature and its impact on your implementation. A key consideration beyond support is the ability of your platform to provide information on sibling CPUs to the SBC . You configure hyperthreading after installation.

VLAN Support

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.

Refer to the *ACLI Configuration Guide* for instructions on configuring VLANs on the SBC. Note that when you configure a VLAN, the SBC requires VLAN tags to be included in the packets delivered to and from the VM.

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

Hyperthreading and CPU Affinity

Hyperthreading can provide increased workload efficiencies in specific configurations, but poorly designed configurations can just as easily impact performance of the SBC.

Due to the polling operation of DPDK, using hyper-threaded cores can significantly degrade the SBC's packet processing performance. Oracle recommends you disable hyper-threading on 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. Learn how to configure CPU affinity via your hypervisor documentation.

To use hyper-threading with SBC, it's important that the hypervisor passes a valid CPU map to the VM, so that SBC has sufficient information to avoid any potential contention from using hyper-threaded sibling cores for realtime critical processes.

In summary:

- Configurations that have hyperthreading disabled on the host are supported in all cases same for both bare metal and virtual hosts.
- Configurations that have hyperthreading enabled on the host are supported if the hypervisor provides correct CPU sibling maps to the guest. This also requires that you enable the use-sibling-core-datapath parameter.
- 3. Configurations that have hyperthreading enabled on the host but do not report sibling maps to the guest are unsupported unless CPU cores are manually pinned at the hypervisor to avoid sibling contention. This also requires that you enable the use-sibling-core-datapath parameter.

You can verify and troubleshoot the SBC CPU assignments using, for example, the **show datapath-config** command and understanding the following guidelines:

- The SBC displays sibling CPUs in lower-case letters:
 - A signaling core with signaling sibling appears as "Ss".
 - Cores other than signaling appear as the core type with no sibling, such as "Dn".
- The SBC displays a verify-config ERROR if there is an error with CPU assignment, including improperly configured hyper-threaded sibling CPUs.


Host Hypervisor CPU Affinity (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.





Oracle recommends you configure CPU affinities on the hypervisor to ensure mapping from only one virtual CPU to each physical CPU core. Learn how to configure CPU affinity and pin CPUs from your hypervisor documentation.



The SBC relies on multiple queues on virtual NICs to scale session capacity. Multiple queues enable the SBC to scale through multiple forwarding cores. This configuration is platform dependent: physical NIC, Hypervisor, virtual NIC, and vSwitch.

Configuration Overview

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

During VM installation, you can configure vSBC boot parameters, including:

- IP address
- Host name

During VM installation, the SBC sets default functionality, assigning cores to signaling and media forwarding. If you need DoS and/or transcoding functionality, you configure the applicable cores after installation. Applicable performance tuning configuration after deployment includes:

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



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 SBC configuration procedures after installation, including:

- Setting passwords
- Setup product
- Setup entitlements
- Assign Cores
- Enable hyperthreading for forwarding, DoS and transcoding cores
- Service configuration

VLAN Support

Refer to the *ACLI Configuration Guide* for instructions on configuring VLANs on the SBC. Note that when you configure a VLAN, the SBC 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.



Virtual Machine Platforms

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

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

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

Create and Deploy on KVM

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

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

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

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 network=medial \
    --network network=media2 \
```

--import \
--cpu host

Note:

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

--name

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

--description

Describe this virtual machine.

--os-type

Specify the operating system type.

--os-variant

Optimize the configuration for a specific operating system.

--ram

Allocate a specific amount of RAM to this virtual machine.

--vcpus

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

--disk

Specify the path to the disk image.

--network

Connect the virtual machine to a host network.

--import

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

--cpu

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

See man virt-install for more information.



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

To maintain performance when using PV mode, you must pre-allocate cores to the Open vSwitch (OVS) for each forwarding core in the VM. For example, if the VM requires four forwarding cores, then four CPU cores should be allocated to the OVS.



Create and Deploy on VMware®

You can deploy Oracle Communications Session Delivery products on ESXI hypervisors running VMware 6. For general deployment tasks, from hardware installation to VM resource management, see the VMware documentation.

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.



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.

- From the VMWare ESXi Host page, click Create/Register VM.
- On the Select Creation Type screen, select Deploy a virtual machine from an OVF or OVA file and click Next.
- On the Select OVF and VMDK Files screen:
 - a. Enter a name for the virtual machine.
 - b. Browse to a target .ova file.
 - c. Click Next.
- 4. On the Select Storage screen, click Next.
- 5. On the License Agreements screen, click I agree and click Next.
- On the Deployment Options screen:
 - a. Map the networks used in this template to networks in your inventory.

The SBC 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:

- i. wancom0
- ii. wancom1
- iii. wancom2
- iv. spare
- v. s0p0
- vi. s1p0
- vii. s0p1
- viii. s1p1

If your hypervisor randomly allocates addresses for network interfaces, the interface order at the hypervisor does not necessarily match that of the SBC. 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. By default, all interfaces use the recommended VmxNet3 driver.

- b. For Disk provisioning, select Thin or Thick.
- c. Leave the box checked to power on automatically.
- d. Click Next.
- On the Additional Settings screen:
 - a. Click **Bootloader parameters** to expand the settings.
 - **b.** Enter the desired boot parameters.



With S-CZ8.4.0p3 or later, you can enter the SHA1 hash of the password for each account (user, admin, and li-admin) into its respective field.

- c. Click Next.
- On the Ready to Complete screen, review the selections, make edits if necessary, and click Finish.

Optimize for VMware

Change these VMware settings to optimize your SBC for performance.

Set VM Compatibility

If you upgrade your hypervisor, you need to manually upgrade the VM compatibility of each previously deployed SBC to match the version of your host.

- Shutdown the VM.
- 2. Select Actions, and then Upgrade VM Compatibility.
- 3. Select the ESXi version that matches the version of the host.
- 4. Click Upgrade.

CPU Pinning

By specifying the CPU affinity settings for each VM, you create VM-to-processor placement constraints. Note that CPU affinity is different than VM-to-VM or VM-to-Host affinity rules.



CPU pinning does not work where vMotion determines VM placement in VMware clusters.

1. Determine the NUMA node assicated with the NIC card used for media traffic.

esxcfg-info -w

Select Settings, and then CPU.



- 3. Set Cores per Socket to the number of CPUs to assign to the VM.
- Set Reservation to the number of CPUs multiplied by the clock speed.
 For example, if the clock speed is 2095 MHz and 8 CPUs are assigned to the VM, reserve 16760 MHz.



Use **show platform cpu** on your SBC to display the CPU clock speed.

- If your hardware supports IOMMU, set IOMMU to Expose IOMMU to the guest OS.
- 6. Under Sheduling Affinity, enter the physical CPU numbers assiciated with the identified NUMA node as comma-seperated numbers.
 - On a hyperthreaded host, use the even numbered CPUs for your SBC VM.
 - Make sure the assigned core numbers, including sibling cores, are not used by other VMs.

You can do this by explicitly assigning the remaining CPUs to other VMs.

For example, if you have 96 available CPUs available, your entry might look like this:

80,82,84,86,88,90,92,94

Memory Pinning

Memory pinning reservations specify the guaranteed minimum allocation for a VM.

- Select Settings, and then Memory.
- Set Reservations to the maximum allocated amount.By default the system reserves only 25% of the allocated amount.

Virtual NIC Settings

By default, all interfaces use the recommended VmxNet3 driver.

Virtual NIC Settings for Media Interfaces

ESXi uses one transmit thread regardless of the number of vNICs associated with a VM. To achieve significant parallelism for high I/O workloads, you can configure a VM to use one transmit thread per vNIC. Enabling this increases network throughput at the cost of higher CPU usage.

- Select your VM, right-click, and then select Settings, and then VM Options, and then Advanced.
- 2. Select Edit configuration and then Add parameter.
- 3. Add a parameter (or edit an existing parameter) called ethernetX.ctxPerDev (where the capital X is the number of the vNIC) and set its value to 1.
 For example, if vNICs 4 and 5 are used as media ports, add the parameters ethernet4.ctxPerDev and ethernet5.ctxPerDev and set each value to 1.
- 4. Click Save.

Latency Sensitivity

For VMs in paravirtualized mode, the **Latency Sensitivity** parameter should be set to **Normal**.



Managing Snapshots

You can use snapshots to preserve the state and data of your SBC prior to upgrading. This allows you to revert back to the previous condition if the upgrade encounters a problem. See VMware's documentation for Using Snapshots to Manage Virtual Machines

Create and Deploy on Hyper-V®

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

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

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


Before You Begin:

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

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

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





- 1. Start the Hyper-V Manager.
- 2. Start the Virtual Switch Manager from the Actions dialog.

Hyper-V displays the Virtual Switch Manager dialog.

3. Click the Create Virtual Switch button.

Hyper-V modifies the **Virtual Switch Manager** dialog, presenting fields within which you specify your new switch.

- 4. Add virtual networks for each management and media interface. Set the following on the **Create Virtual Switch** dialog for each switch you create:
 - Virtual switch Name
 - Uncheck Allow management operating system to share this network adapter
 - The switch for wancom1, wancom2 can be **internal** or **external**. High availability via external (eq., separate hypervisor platforms) is preferred.
 - All other switches must be external.

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

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



h. s1p1

Click New, Virtual Machine

Hyper-V displays the introductory page of the New Virtual Machine Wizard.

Click Next.

Hyper-V advances through the **New Virtual Machine Wizard** pages each time you click **Next**. The wizard allows you to go back to the Previous page, Cancel the wizard, or Finish the wizard with the respective buttons. Your procedure through the wizard may vary, depending on your infrastructure and intent.

- 7. Enter or select at least the following as you progress through the wizard.
 - Type a name for your VM in the Name field.
 - **b.** Select **Generation 1** as your machine type.
 - c. Assign the desired memory.
 - d. Click Next, skipping the Configure Networking dialog. You add networks later in the process.
 - e. Connect to the Virtual Hard Disk you downloaded by selecting the **Use an existing virtual hard disk** radio button and browsing to your .vhd file.
 - f. Finish

The Hyper-V Manager returns to the main dialog, displaying your new machine in the Virtual Machine list.

8. Right click the VM.

The Hyper-V Manager displays a pop-up menu.

Click Settings

The Hyper-V Manager displays the **Settings** dialog for your Virtual Machine displaying the **Add Hardware** controls in the right-hand pane.

10. Select Legacy Network Adapter and click the Add button.

The Hyper-V Manager displays the **Legacy Network Adapter** dialog.

11. Select wancom0 you configured for your Virtual Machine from the drop-down listbox and click the Apply, then the OK buttons. Repeat this step for wancom1 and wancom2 if you are using these interfaces.

The Hyper-V Manager returns to the initial Settings dialog and adds this adapter to your machine's component list in the left-side pane. Only configure wancom0, wancom1 and wancom2 as **Legacy Network Adapters**.

12. Select Network Adapter and click the Add button.


The Hyper-V Manager displays the **Network Adapter** dialog.

- 13. Select the first adapter after wancom0 that you configured for your Virtual Machine from the drop-down listbox and click the **Apply**, then the **OK** buttons.
- 14. Repeat the previous step for the rest of your adapters, referring to the order described above.
- 15. Re-Open the **Settings** ... dialog.
- For each Network Adapter, click the + sign beneath it to display the Hardware Acceleration Advancedlinks.
- For each Network Adapter's, Hardware Acceleration settings, uncheck the Enable virtual machine queue checkbox. Apply and OK these changes.



18. For each media interface's Network Adapter, Advanced settings, check the Enable MAC address spoofing checkbox. Apply and OK these changes.

Figure 5-1 VM Settings on Hyper-V



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



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



OpenStack Heat Template

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

Extract the Environment File

After downloading the compressed TAR file from Oracle, extract the contents on the machine from which you will deploy virtual machines.

1. Extract the compressed HOT file bundled with your software package.

```
tar xzf <filename> HOT.tar.gz
```

Extract the Newton tar file if your OpenStack is running Newton; extract the Pike tar file if your OpenStack is running Pike or newer.



For OpenStack versions newer than Pike, see the VNF Metadata and Userdata Example appendix.

```
tar xf <filename>_HOT_newton.tar
```

This creates a local directory that contains the environment files.

- Locate the environment file in the properties directory.
 - For HA environments: properties/sdHaParams.yaml
 - For standalone environments: properties/sdStandaloneParams.yaml

Select Product and Entitlements

Both the product and its entitlements are defined in the entitlement file. Entitlement files are located in the entitlements directory.

Setup Product Description	Entitlement File Name
Session Border Controller	sbc.yaml
Peering Session Border Controller	perring_sbc.yaml
Enterprise Session Border Controller	esbc.yaml
Session Router - Session Stateful	<pre>sr_session_stateful.yaml</pre>
Session Router - Transaction Stateful	<pre>sr_transaction_stateful.yaml</pre>
Subscriber-Aware Load Balancer	slb.yaml

Entitlement files are passed to the openstack command when deploying a virtual machine. For complete instructions on deploying a Heat template, see the README.



Set Parameters in Environment File

The following parameters can be configured in the environment file. Note that some parameters are only available in the HA environment file and not in the standalone environment file.

- primaryName—Name of the instance as displayed in the OpenStack GUI. The value gets passed in the bootparams as the Target Name.
- secondaryName—(HA only) Name of the instance as displayed in the OpenStack GUI. The value gets passed in the bootparams as the Target Name.
- highAvailability—Enable or disable HA for this template. Always set to true in the HA environment file and false in the standalone environment file.
- enableRestInterface—Enable or disable the REST API interface. By default this is set to false.

If you set enableRestInterface to true, the SBC generates a self-signed certificate to enable the REST interface. If your REST client requires a specific TLS version or key size, you may edit the files in the xmlconfiq directory to change the properies of the temporary self-signed certificate.



WARNING:

Replace the self-signed certificate with your own CA-signed certificate before moving the SBC into a production environment.

- flavor—The name or ID of the OpenStack flavor.
- image—The name or ID of the previously uploaded SBC image.
- availabilityZone—Specify the availability zone where the SBC will be placed.
- securityGroup—Specify the security group.



Note:

The default security group drops incoming traffic, so either change the parameters of the default security group or create a new security group that allows incoming traffic.

- affinityPolicy—(HA only) Set the affinity policy for an HA pair. The default value anti-affinity is recommended.
- userPass—The SHA1 hash of the user passphrase.
- adminPass—The SHA1 hash of the admin passphrase.
- diskPartitions—Specify the percentage of disk space that will be allocated for each partition.
- licenseKeys—The license keys to be used when the SBC is created.
- applyBaseConfiguration—Enable or disable the base configuration which is suitable for minimal Standalone or HA-pair functionality.
- configuration—If applyBaseConfiguration is set to true, specify the input parameters for the base configuration. Subparameters include:



- dosCores—Specify the number of CPU cores dedicated for denial-of-service protection.
- forwardingCores—Specify the number of CPU cores dedicated for forwarding frames.
- transcodingCores—Specify the number of CPU cores dedicated for transcoding media.
- ntpServer1—Specify the IP address of an NTP server to use for time synchronization.
- ntpServer2—Specify the IP address of an NTP server to use for time synchronization.
- snmpCommunityName—Specify the name of the SNMPv2 community to use for SNMP management.
- snmplpAddress—Specify the IP address to add to the SNMPv2 community for SNMP management.
- wancom0VLAN—(Only available on Pike and newer) Specify the bootparameter VLAN value for the wancom0 interface.
- networks—Specify the name of the networks to attach to each network interface.
- portSecurityPolicy—Specify the port security policy currently in use for each network.
- primary:fixed_ips—Specify the IP addresses for the network interfaces on the primary HA pair. The wancom0 address gets passed in the bootparams as IP Address.
- secondary:fixed_ip—(HA only) Specify the IP addresses for the network interfaces on the secondary HA pair. The wancom0 address gets passed in the bootparams as IP Address.
- virtualIPs—(HA only) Specify the virtual IP addresses for HA media interfaces.
- virtualMACs—(HA only) Specify the virtual MAC addresses for HA media interfaces.
- vnicBinding—Specify the virtual NIC binding type for each media interface.

See the comments in the environment file for available values.

See the README for complete instructions on deploying a Heat template.



6

Public Cloud Platforms

Oracle distributes machine templates, each containing a virtual disk image and default configuration for the supported profile of each public platform. Public Cloud platform support is dependent on your Oracle product version. Refer to your Oracle Communications Session Border Controller (SBC) *Release Notes* to confirm the public clouds supported and important detail on that software version's support.

This section addresses requirements associated with running the SBC as public cloud instances. It also provides basic instructions on deploying machine instances.

Public Cloud providers maintain extensive product documentation. You must use those vendors' documentation for specifications, requirements, caveats, known issues, deployment details, and operation detail prior to deploying the SBC.

Once you have installed on a public cloud, the SBC provides the same or equivalent functions that allow access to and configuration of the virtual SBC. Access functions include SSH, Console and external management applications. Configuration functions include product setup, entitlements setup and ACLI engine.



After installation, you must setup entitlements. If you later change machine shape or size, your new deployment may not accommodate the original entitlements, such as capacity entitlements. When you complete machine size or shape changes, run the **show platform limits** command and note the supported capacity limits. After changing virtual machine size or shape, change any entitlements that conflict with you new platform's support and reboot your virtual machine.

Deploying the SBC on Cloud Infrastructures in Standalone Mode

Deploying the SBC in standalone mode on public clouds consists of deploying a single SBC. Standalone deployments are significantly less complex to configure and maintain. Although some object types are not required for standalone mode, you can use them for other purposes, if desired.

Differences between standalone and HA deployment include:

- EC2
 - The IAM role is not required for the SBC
 - Only a single management interface is required
 - No Elastic IP addressing is required
 - No Place Group is required
- OCI
 - No Instance Principal Authorization is required

- No Dynamic Group is required
- Only a single management interface is required
- No Secondary Private addressing is required
- Azure
 - The IAM Network Contributor role is not required for the SBC
 - No Fault Domain is required
 - No Availability Set is required
 - Only a single management interface is required
 - No Elastic IP addressing is required

Although the procedures in this document include steps that are used to support HA, you use the same procedures without the HA steps to deploy the SBC in standalone mode. Procedure documentation indicates when a step applies to HA deployments.

Deploying the SBC on Cloud Infrastructures in HA Mode

The Oracle Communications Session Border Controller (SBC) supports High Availability (HA) deployments on public clouds using the redundancy mechanisms native to those clouds.

If you are deploying on Oracle Cloud Infrastructure (OCI), you can enable the SBC's native, more efficient GARP-based HA. To use this method:

- Open a Support Ticket for your OCI tenancy.
- At the direction of your Oracle Support representative, disable the OCIHAClient SPL plugin.

```
ORACLE (configure terminal
ORACLE (configure) # system spl-config
ORACLE (spl-config) # sel
ORACLE (spl-config) # plugins
ORACLE (spl-plugins) # name OCIHAClient.1.0.spl
ORACLE (spl-plugins) # state disabled
ORACLE (spl-plugins) # done
```

3. Save and activate your configuration.

On all other cloud infrastructures (or on OCI with GARP-based HA not enabled), you configure the cloud to recognize the SBC. The REST client on the SBC subsequently makes requests to the cloud's Software Defined Networking (SDN) controller for authentication and virtual IP address (VIP) management. While HA configuration across all SBC platforms is similar, public cloud HA configuration fundamentally does not require configuring virtual MAC addresses. This feature supports only IPv4 addressing. The SBC includes a REST client to configure the cloud's SDN controller. The local REST client supports both HTTP and HTTPS, using the former for metadata requests and the latter for other cloud management requests. The SBC does not support using both GARP-based HA and REST-based HA simultaneously.

Vendors manage public clouds using SDN. The SDN controller owns all networking aspects including vNICs, IP addresses, MAC addresses, and so forth. Without the knowledge of the SDN controller, IP addresses cannot be assigned or moved. As a result, the network either drops or ignores GARP traffic. The absence of GARP invalidates the use of HA by the SBC in these networks, therefore requiring alternate HA functionality on the SBC.

The SBC recognizes when it is deployed on these clouds. When it needs to failover, instead of issuing GARP traffic to invoke the transfer of VIPs from one node to another, it uses the cloud's REST API to reconfigure virtual IP addressing.

Cloud configuration and the use of REST is equivalent across the range of public clouds, with vendors using different terminology for similar functions and objects.



The SBC does not support High Availability (HA) when deployed over Azure.

It is recommended not to configure the boot parameter **IP address** while deploying SBC on public clouds as DHCP gets disabled and no DNS Server is available that results in "could not resolve the host" error while switchover happens.

Key SBC Configuration and Operation Detail in HA Mode

When the SBC is deployed in an HA configuration in a public cloud, the system's configuration and operation is predominantly the same as in an on-premise HA configuration. The only operational difference is its behavior when going active. Because the SBC knows it is deployed on a public cloud, it automatically replaces its GARP procedures when it goes active with REST calls to fetch VIPs. Going active includes first startup of the primary, as well as the standby taking over because of an HA event.

HA Operation

When an SBC goes into active state on a public cloud, its REST client requests VIPs. If it does not receive the addressing, it re-issues these requests after 5 seconds, then 10, 20, 40, 80 and finally 160 seconds. If these request fail, it attempts to acquire VIPs every 300 seconds until it succeeds. Upon success, the SBC suspends these re-requests and begins to send address verification requests every 300 seconds. These requests verify that the SDN continues to associate this SBC with these VIPs.

When moving from Active to any other state, the SBC gracefully abandons any outstanding REST client operations.

Configuration Considerations

Key SBC configuration considerations across all cloud environments include:

- Do not configure virtual MAC addresses
- Configure VIPs via the cloud's console as secondary private IPs on the media vNICs
- Obtain wancom0 management address via DHCP
- Obtain the wancom0 default gateway via DHCP
- Obtain the cloud's name servers via DHCP, allowing the local DNS resolver to cache cloud infrastructure addressing

Regarding the use of DHCP to obtain addressing, this means that those subnets must have DHCP enabled and DNS name server IPs configured.

Authentication

The public cloud's API requires that the client authenticate with the cloud to successfully invoke the API. Although each cloud has differences between their authentication mechanisms, they



are typically categorized as either Provisioning or Automatic. The SBC uses Automatic authentication.

With Automatic authentication, the SDN assigns an SBC instance with an appropriate role, providing credentials through its instance metadata. The instance does not need to be provisioned with any credentials. These credentials may be temporary and change periodically. As a result, the SBC instance does not cache its credentials, obtaining the latest credentials when invoking the API.

DNS Resolver


Because the REST API resolves hostnames into IP addresses, SBC needs an active DNS resolver. This ensures that it can forward properly. As part of this feature, the SBC enables its DNS resolver when deployed on public clouds. The SBC then obtains the name server IPs for the DNS resolution through DHCP. The DHCP client obtains the name server IPs and populates them in the *letc/resolv.conf* file. The SBC uses the nameservers for DNS resolution when invoking REST API.

Note: The SBC must obtain the IP address for the wancom0 interface through DHCP which will provide the name server IPs. Without DHCP, the SBC has no mechanism to manually configure the name server IPs.

IP Addressing

The following diagram portrays an IP addressing example applicable to an HA deployment in a public cloud. Note that the external endpoints are SIP peers reachable through internet/public IPs.





The minimum number of subnets is four:

- 1 management (wancom0)
- 1 HA subnet (wancom1)
- 2 media subnets (s0p0 and s1p0)

Additional subnets may include:

- 1 additional HA subnet (wancom2)
- 6 additional media subnets, a total of 8 media subnets

Cloud categories of typical IP addressing includes:

- Private IP—Each vNIC has a primary private IP, and you can add and remove secondary private IPs. The primary private IP address on an instance does not change during the instance's lifetime and cannot be removed from the instance.
- Secondary Private IP—Each VNIC can be assigned additional IP addresses (apart from the primary private IP), from the same subnet, called secondary private IPs.
- Public IP—Each VNIC with a private IP can optionally be assigned with a Public IP. Public IPs can be either:
 - Ephemeral Public IP—A temporary public IP and exists only for the lifetime of the instance.

 Reserved Public IP (Elastic IP in AWS)—Persistent and exists beyond the lifetime of the instance it's assigned to. You can unassign it and then reassign it to another instance whenever you like.

As you configure IP addressing, note that both OCI and AWS use the terminology, **Primary Private IP**, which maps to the SBC **primary-utility-address** and **Secondary Private IP** which maps to the SBC **secondary-utility-address**. Important considerations include:

- A management subnet is typically a public subnet. Each instance that requires
 management access from outside the cloud needs a public IP assigned. Oracle
 recommends you use a Reserved Public IP instead of Ephemeral Public IP. HA subnets
 are private subnets.
- Media subnets can be either public or private subnets depending on where the SBC peers are located, as follows:
 - If the SBC peers reside in the cloud on the same subnet, you need a private subnet.
 - If the SBC peers are reachable through Internet, you need a public subnet.
- For each media subnet the SBC is attached to, you need one secondary private IP to act as VIP. Additionally, for each media subnet that requires Internet access, the secondary private IP must be attached to a Reserved Public IP.

RTC Support

This feature is RTC supported. If you make any configuration changes that affect HA operation, the SBC immediately issues request for new VIP addressing via its REST client.

Cloud-Based HA Report Data

The SBC logs all events related to VIP failover for each VIP including failure, number of attempts, nature of response received from the REST API, and the end result. You can observe the system's REST requests and resultant processing in the HA operation log.

The SBC logs all REST client events into the **log.cloudha** file. The SBC generates alarms for all REST API failures and clears all alarms as soon as the respective REST API failure clears.

SBC REST API alarms include:

- Alarm code APP_ALARM_CLOUD_HA_USER_ERROR (#327732), includes the following CRITICAL alarms:
 - IAM role does not exists
 - Invalid credentials
 - User has insufficient permissions
 - Your account is currently blocked
- Alarm code APP_ALARM_CLOUD_HA_CURL_ERROR (#327733), includes the following alarms. These alarms are of type CRITICAL, unless specified otherwise:
 - Cloud not resolve host
 - cURL timedout to invoke REST API (type MAJOR)
 - cURL failed to invoke REST API
- Alarm code APP_ALARM_CLOUD_HA_SERVER_ERROR (#327734), includes the following CRITICAL alarms:
 - The server is overloaded and can't handle the request



- The request has failed due to a temporary failure of the server
- Alarm code APP_ALARM_CLOUD_HA_API_ERROR (#327735), includes the following alarms. These alarms are of the type CRITICAL, unless specified otherwise:
 - Invalid secondary private IP (type MAJOR)
 - A parameter specified in a request is not valid, is unsupported, or cannot be used
 - Failed to switchover VIP
 - Invalid usage of the HA script

Cloud-Specific HA Deployment Considerations

Deployment over individual cloud environments pose specific authentication and configuration requirements. In addition, each environment uses region types that you must adhere to so the SBC can perform HA functions.

HA over OCI Overview

OCI uses **Instance Principal Authorization** to allow the instances to access services. The following steps summarize the process flow for setting up and using instances as principals. Upon completing these steps, the SBC instance can then obtain a temporary certificate to authenticate itself while invoking the API.

- 1. Create a **Dynamic Group**. In the dynamic group definition, you provide the matching rules to specify the instances you want to allow to make API calls for services.
- 2. Create a policy granting permissions to the dynamic group to access services.

As you deploy, follow these guidelines:

- Create both SBC instances in the same Availability Domain
- Oracle recommends that you create SBC instances in separate Fault Domains.

As you configure, follow these guidelines:

- Do not configure and use more that 4 secondary private IP addresses per HA deployment.
 More than 4 IPs causes HA failover to take too long.
- On the primary SBC instance, configure Secondary Private IPs (to be used as SBC virtual IPs) through the OCI console. Do not use the SBC ACLI to configure a **sec-utility-addr**.
- When required, map your Secondary Private IPs to Reserved Public IPs.

HA over AWS Overview

AWS uses **Identity and Access Management** (IAM) roles to provide instances access to the infrastructure services. Configure an IAM role with required policies and associate the IAM role with SBC instances during creation. The instances can then obtain the credentials through the metadata and authenticate itself while invoking the API.

As you deploy, follow these guidelines:

- Create both SBC instances in the same Availability Zone.
- Oracle recommends that you use Place Groups of type Spread for launching both SBC instances.

As you configure, follow these guidelines:



- On the primary instance of SBC, and through the AWS console, configure Secondary Private IPs to be used as SBC virtual IPs.
- 2. When required, map Secondary Private IP addressing with Elastic IP addressing.

AWS uses its **Access key ID** and **Secret access key** as security credentials. Since these credentials change periodically, the SBC does not cache the information. Instead, the SBC always retrieves and uses the latest information from the metadata. In addition, the SBC retries the API by refreshing the latest security credentials if it receives any error response indicating the authentication failed.

Create and Deploy SBC Instances on EC2

You can deploy the Oracle Communications Session Border Controller (SBC) on Amazon's Elastic Computing (EC2) infrastructure in either standalone or High Availability (HA) mode.

When deployed on this platform, you configure and operate the SBC as you would on any other platform. You deploy the SBC to use the environment's IP infrastructure, including the private and public addressing scheme, and its translation functions to protect the EC2 management environment from direct exposure to the SBC service delivery environment. These deployments also use EC2's DHCP to establish consistent, compliant management addressing and flat networking across both management and service traffic.

The SBC does not support using AWS Direct Connect for HA on management interfaces.

Oracle recommends you use AWS enhanced networking for better network I/O throughput. Refer to your SBC Release Notes for your version's supported machine shapes. You can also refer to AWS resources, including the Enable enhanced networking with the Intel 82599 VF interface on Linux instances web page for key information about enhanced networking, including the EC2 machine shapes that support it.

Find release-specific requirements and information on this feature in your *SBC Release Notes*. Find configuration information in the *SBC ACLI Configuration Guide*. If desired, request information to refine your virtual machine resources configuration to desired performance from Oracle Support or Sales.

Prerequisites to EC2 Deployment

Prerequisites to this deployment procedure include:

- You have identified and are deploying to the correct AWS Region. This is typically a
 default component of your EC2 Account.
- You have identified and are deploying to the correct AWS Availability Zone. By deploying 2 (HA) instances during deployment at the same time, you are ensuring that both instances reside in the same Availability Zone.
- An Amazon Virtual Private Cloud (VPC) is configured.
- A security policy is configured.
- You have determined the number of management and media interfaces you want for each instance.
- All subnets are configured. Each SBC management and media interface requires its own unique subnet.

When deploying within an AWS VPC, AWS provides a predefined DNS Server IP. This means that you must use caution when determining your service networks. AWS predetermines this address as the VPC CIDR range +2.

For example, if AWS provides you with a VPC CIDR Range of 10.8.0.0/16, it also provides the DNS IP as 10.8.0.2 using DHCP. If you have configured any media or HA subnet such that it overlaps this network, such as 10.8.0.0/22, the SBC sends DNS requests out that interface, which eventually fail.

Your EC2 workspace may present dialogs and fields that differ from this procedure. For full information on deploying EC2 instances, see the Amazon EC2 documentation.

Generate an EC2 AMI from the vSBC Image

You perform this procedure to convert the Oracle Communications Session Border Controller (SBC) image provided by Oracle into an AMI, from which you can create vSBC machines.

This procedure requires that you, from the EC2 system, create a Linux machine, attach a new disk to it, and put the SBC image on your machine. Next, use the Linux command line to convert the image from qemu to raw format, then perform a data definition (dd) procedure to write the disk image. Finally, you create a snapshot of the disk and convert the snapshot to an AMI image. You can create vSBCs from this AMI.

- 1. Launch an instance of the Amazon Linux 2 (HVM) with a Public IP for ssh access. You must create this Linux instance in the same Availability Zone as your second disk volums.
- 2. Attach a second disk volume of 20GB to the VM instance, as a known device (e.g. /dev/xvdb). Create this disk using volume type **General Purpose SSD** (gp2).
- Download your SBC image via Oracle support. The correct image name is appended with img-vm_kvm.tgz.
- 4. Copy the KVM release image to the target VM.

```
scp -i "Oracle.pem" nnSCZ830-img-vm kvm.tgz ec2-user@public ip addr:
```

5. ssh into the target VM.

```
ssh -i "Oracle.pem" ec2-user@public ip addr
```

6. Extract the qcow2 disk imagie from the release package. The command line below uses version 8.3.0 as an example.

```
tar xvfz nnSCZ830-img-vm_kvm.tgz
```

7. Install the KVM image conversion utility.

```
sudo yum install qemu-img
```

8. Convert the image to raw format.

```
qemu-img convert -p nnSCZ830-img-vm kvm.qcow2 nnSCZ830-img-vm kvm.raw
```

9. When you attach the disk to the VM, choose the disk symbol, for example **sdf**. You can check this by navigating to Image/Volumes/Attachment information.

```
Attachment information i-06715fac7647142fc (SQA-OL6-ANVIL-1):/dev/sdf (attached)
```



Verify the attachment using the linux list block devices command, Isblk. Example output is shown below.

```
host:/dev$ lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
xvda 202:0 0 200G 0 disk
Lxvdal 202:1 0 15G 0 part /
xvdf 202:80 0 20G 0 disk
```

11. Write the raw image into the 20GB volume.

```
sudo dd if=nnSCZ830-img-vm kvm.raw of=/dev/xvdb
```

You may or may not be able to use the status=progress flag within this command. This is dependent on the version of the dd command. In addition, note that using the status=progress flag causes the process to take longer.

- 12. From the EC2/Elastic Block Store/Volumes tab, right-click the 20GB volume and click **Detach Volume** from the VM.
- **13.** From the same tab, right click the 20GB volume again, click **Create Snapshot** and wait for the snapshot to complete.
- 14. From the EC2/Elastic Block Store/Snapshots tab, right-click the snapshot, and click **Create Image** to generate an AMI.

The Create Image from EBS Snapshot dialog oprovides for multiple configurations, including:

- Name your Image.
- Select x86_64 from the Architecture drop-down.
- Select Use default from the RAM disk ID drop-down.
- Select **Hardware-assisted virtualization** from the Virtualization type drop-down.
- Select Use default from the Kernel ID drop-down.
- Select General Purpose SSD (gp2) from the Volume Type drop-down.

You can use this AMI to launch new VM instances. Follow the procedures to format the hard disk if you create instances with a disk volume greater than 20GB.

EC2 Deployment Procedure

Deploying the SBC on EC2 includes the following high-level steps:

- Launch your Instances on AWS—This is the main instance configuration procedure. It
 includes a multi-dialog wizard that presents configuration options in the preferred
 sequence. The result of this wizard is an installed, operational SBC or two HA SBC
 instances with no networking.
- 2. Configure the Network Interfaces for Your Instances—This is a preparatory task, creating interfaces that you attach to instances in the next procedure.
- Attach the Network Interfaces to Your Instances—This assigns interfaces to instances.
 Check your interface assignments using the SBC ACLI interface-mapping commands, shown below, after SBC startup and correct interface assignment, if necessary. See the

Oracle® Communications Session Border Controller Platform Preparation and Installation Guide for further instructions on using these commands.

- Apply your SBC Configuration—This is an SBC ACLI configuration procedure.
- 5. During ACLI configuration, configure the first private address provided by EC2 for media interfaces as the pri-utility-address, and the secondary private address as the interface's ip-address when you configure your primary HA SBC.
- 6. During ACLI configuration, you only need the first private address provided by EC2 for the secondary HA SBC media interface addresses. Configure that address as each applicable media interface's sec-utility-address.

Deploying the SBC for HA, adds the following high-level considerations:

- Configure an Identity and Access Management (IAM) Role for the SBCs—This is a preparatory task.
- Place both SBCs in the same Availability Zone.
- A Place Group, of the type Spread, must be available, within which you place both SBCs.
- You must assign a public IP to the wancom0 management interface.
- Configure Secondary Private IPs (Virtual IPs) for all Media Interfaces to create virtual IPs for use during HA switchovers.
- You must map the Secondary addresses used for Virtual IPs to Elastic IP addressing.
- The wancom0 management subnets must be public to allow access from outside the cloud. You can meet this requirement by allowing an auto-assigned Public IP or by configuring it with an Elastic IP.

For both HA and standalone deployments, all Media interfaces addresses that must be reachable through the internet must reside on public subnets; all others can reside on private subnets. In addition, you can create additional IP addresses for an interface, allowing for different addresses on **steering-pools**, HIP, and other objects.

Your EC2 workspace may present dialogs and fields that differ from this procedure. For full information on deploying EC2 instances, see the Amazon EC2 documentation.

Configure an IAM Role


Create an IAM policy and role for high availability SBC instances. (This is not required for standalone SBCs.)

- Create an HA policy.
 - a. Navigate to Services, and then IAM, and then Policies, and then Create policy.
 - b. Select the JSON tab.
 - Paste the following JSON into the JSON editor.




```
"ec2:AssociateAddress",
    "ec2:DescribeSecurityGroups",
    "ec2:DescribeInstanceStatus"
],
    "Effect": "Allow",
    "Resource": "*",
    "Sid": "VisualEditor0"
}
],
    "Version": "2012-10-17"
}
```

d. Click Review Policy.



- e. Enter a name and description for this policy.
- f. Click Create policy.
- 2. Create an HA role.
 - a. Navigate to Services, and then IAM, and then Roles, and then Create role.
 - b. Select AWS service, and then EC2, and then Next: Permissions.
 - c. Select the previously created policy and click **Next: Tags**.
 - d. Add metadata tags, if desired, and click Next: Review.
 - e. Enter a name and click Create role.




Create a NAT Gateway


For customers who want no direct connectivity between the wancom0 interface and the public internet, create a NAT gateway. For details on pricing, see Amazon VPC pricing.

- Create a NAT gateway.
 - Navigate to Services, and then VPC, and then NAT Gateways, and then Create NAT gateway.
 - b. Enter a name.
 - c. Select a subnet.
 - d. Select an Elistic IP.
 - e. Click Create NAT gateway.
- Create a route table.
 - Navigate to Services, and then VPC, and then Route tables, and then Create route tables.
 - b. Enter a name.
 - c. Select a VPC.
 - d. Click Create.
- 3. Add the NAT gateway and private wancom0 subnet to the route table.
 - a. From the Route Tables page, select your route table.
 - b. In the Routes tab, select Edit routes and then Add route.
 - c. Add a route for 0.0.0.0/0 to the target route.





- d. Click Save routes.
- 4. Associate your private wancom0 subnet with the route table.
 - a. From the Route Tables page, select your route table.
 - b. In the Subnet associations tab, select Subnet associations and then Edit subnet associations.
 - Select your wancom0 subnet.




d. Click Save.

Launch Your Instance

- Login to the AWS management console and click the EC2 link to open the EC2 Dashboard.
- 2. Review and confirm your deployment's Region and Availability Zone.
- 3. On the EC2 Dashboard, click Launch Instance.





- 4. Navigate to the My AMIs link to choose the image for your instance, and click Select.
- Chose the desired instance type. See your software version's release notes for tables of supported machine sizing.
- **6.** Click **Next: Configure Instance Details**. The AWS instance deployment sequence displays the **Configure Instance Details** dialog.
- 7. Configure the following instance details; leave the others at their defaults:
 - a. Specify the number of Instances. (Specify 2 for an HA setup.)
 - Select the correct Network for wancom0.
 - c. Select the correct Subnet for wancom0.
 - **d.** Establish a public IP for wancom0, either by using the **Auto assign Public IP** control or by configuring an elastic IP after deployment.
 - **e.** HA only—Check **Placement Group**. Ensure the group is of type **Spread**, and that both SBCs reside in the same group.
 - f. HA only—Select the appropriate IAM role. (Choose the IAM role you configured above.)
 - g. Scroll down to the **Network interfaces** configuration fields.
 - Ensure you are configuring the Device named eth0.
 - Select New network interface from the Network Interface dropdown list for wancom0.
 - j. Select the correct **Subnet** from the dropdown list for wancom0.
 - k. Ensure the Primary IP field is set to Auto-assign.
- 8. Scroll to the bottom of the **Configure Instance Details** dialog and click **Next: Add Storage**.
- 9. Choose your desired SBC storage size in GB. The default storage size is 40GB.
- 10. Click Next: Add Tags.



- 11. Enter any arbitrary name to identify the instance. Ensure the name allows you to uniquely identify this instance during later deployment procedures and operation.
- 12. Click Next: Configure Security Group.
- 13. You can either create a new security group or select an existing security group to set appropriate firewall rules. Refer to EC2 documentation for configuration instructions.
- 14. Click the Review and Launch button. EC2 displays a summary of your instance.
- 15. Review the Instance configuration and click the **Launch** button.
- **16.** From the pop-up screen, select an existing SSH key pair or create a new key pair and check the **acknowledgment** check-box.
 - a. If you create a new key pair, enter a name and click **Download Key Pair**.
 - b. Move the PEM file to your .ssh directory.
 - c. After launching the instance, SSH to the VM with the -i argument.

```
ssh -i .ssh/my new keypair.pem admin@<Public IP address>
```

- 17. Click Launch Instances. EC2 creates your instances.
- 18. Return to the EC2 Dashboard and click the Running Instances link.
- Select your new instances and name them. These names can be the same as your tag names.

Create Network Interfaces for SBC Instances

- From the EC2 Dashboard, click Network Interfaces under Network & Security on the left panel.
- Click Create Network Interfaces.
- Create HA and/or Media interfaces by selecting the appropriate subnet and security group from the popup.

Example configurations on an instance named myHA1 include:

- Wancom-1 interface for myHA1
 - For **Description**, type in a name that you can clearly recognize later.
 - For Subnet, choose the subnet you created for HA management from the dropdown.
 - For Private IP, retain the auto assign setting, based on the following criteria. If you use auto assign, EC2 applies the first available IP from the subnet to that interface. If you need more precise IP management, the custom option is recommend.
 - For Security groups, choose the Security group you created for this management from the drop-down.
- s0p0 media interface for myHA1
 - For **Description**, type in a name that you can clearly recognize later.
 - For **Subnet**, choose the subnet you created for this media interface from the dropdown.
 - For Private IP, retain the auto assign setting.
 - For Security groups, choose the Security group you created for media from the drop-down.



Perform this step for each management and media interface on your instance.

Attach the Network Interfaces to the SBC Instances

- From the EC2 Dashboard, click Running Instances.
- Select your first instance. Ensure that it is highlighted
- 3. Open the Actions drop down and select Networking, Attach Network Interface.
- From the Attach Network Interface pop-up, select your first network interface name.
- 5. Repeat these steps for all network interfaces created above.
- 6. Repeat these steps for all your instances.

Configure Secondary Private IPs (Virtual IPs) for all HA Deployments

This procedure, which creates virtual addressing, applies only to HA deployments. Perform these steps on the Primary instance of the HA pair only.

- 1. From the EC2 Dashboard, click Running Instances.
- From the bottom panel, select Description, Network Interfaces.
- 3. Click one of the media interfaces. Its network interface details appear in a pop-up.
- 4. Click **Interface ID** from the pop-up window. This takes you to the network interface that is mapped to this media interface.
- From the Actions link, click the Manage IP Addresses option. This opens the Manage IP Addresses pop-up.
- Click the Assign new IP option. This assigns a new secondary private IP address to the network-interface selected.
- 7. Click the Yes, Update button.
- 8. Repeat these steps for all the media interfaces on the current instance.

Configure Elastic IP Addressing

This procedure, which makes virtual addresses persistent through HA switchovers, applies only to HA deployments.

- Under Network & Security in the left column, click Elastic IPs.
- 2. Click Allocate Elastic IP Address, and then Allocate, and then Close.
- Select the newly allocated IP address and click Actions, and then Associate Elastic IP Address.
- 4. Click on the text box next to **Instance** and select your instance from the drop-down menu.
- 5. Click Associate.

Set the User and Administrative Passwords on the SBC

These password procedures are required before any further SBC operations. For HA deployments, perform these procedures on both SBCs.

- 1. From the E2C Dashboard, under Instances in the left column, click **Instances** and click the newly created SBC.
- 2. Under the Description tab, note the public hostname and the Instance ID.



- 3. When the virtual machine has finished initializing, SSH to the public hostname. The username is "user" and the initial SSH password is "acme" + the instance ID.
- 4. Set the user password by logging in for the first time.

```
$ ssh user@somewhere.compute-1.amazonaws.com
user@somewhere.compute-1.amazonaws.com's password:
*ALERT*
******************
user password has not been set. Please set password now.
******************
** Only alphabetic (upper or lower case), numeric and punctuation
** characters are allowed in the password.
** Password must be 8 - 64 characters,
** and have 3 of the 4 following character classes:
** - lower case alpha
** - upper case alpha
** - numerals
** - punctuation
******************
Enter New Password:
Confirm New Password:
```

5. Set the administrative password by typing **enable** at the command prompt. The initial enable password is "packet" + the instance ID.

```
> enable
Password:
*ALERT*
*******************
admin password has not been set. Please set password now.
************
** Only alphabetic (upper or lower case), numeric and punctuation
** characters are allowed in the password.
** Password must be 8 - 64 characters,
** and have 3 of the 4 following character classes :
** - lower case alpha
** - upper case alpha
** - numerals
** - punctuation
*****************
Enter New Password:
Confirm New Password:
```

Verify the network interfaces have MAC addresses.
 Use the show interfaces mapping command to verify the network interfaces have MAC addresses.

```
# show interfaces mapping
Interface Mapping Info
```



```
Eth-IF MAC-Addr
                              Label
wancom0 06:DF:71:BA:D8:77
                              #generic
wancom1 06:A6:08:58:92:C9
                              #generic
     06:D4:E6:E8:B8:FB
                              #generic
       06:EA:08:51:4D:DF
                              #generic
s1p0
wancom2 FF:FF:FF:FF:FF
                              #dummy
spare FF:FF:FF:FF:FF
                              #dummy
                              #dummy
s0p1
       FF:FF:FF:FF:FF
s1p1
       FF:FF:FF:FF:FF
                              #dummy
s0p2 FF:FF:FF:FF:FF
                              #dummy
s1p2
    FF:FF:FF:FF:FF
                              #dummy
s0p3
    FF:FF:FF:FF:FF
                              #dummy
s1p3
       FF:FF:FF:FF:FF
                              #dummy
```

Execute the **interfaces-mapping**, **swap** command, if necessary, to correct any issues with your interface to MAC address mapping.

7. Reboot the virtual machine.

```
# reboot
```

Refer to the *Oracle® Communications Session Border Controller Configuration Guide* after you have completed this deployment for administrative and service configuration, including product setup, entitlement setup and HA configuration.

Deploy AWS AMI with Terraform

Use terraform to deploy the SBC AMI on AWS.

Prerequisites

- Create an IAM policy for terraform to create an S3 stack.
 - a. Navigate to Services, and then IAM, and then Policies, and then Create policy.
 - b. Select the JSON tab.
 - c. Paste the following JSON into the JSON editor.

```
"s3:ListBucketVersions",
                "s3:GetBucketLogging",
                "dynamodb:ListTagsOfResource",
                "s3:CreateBucket",
                "s3:ListBucket",
                "s3:GetAccelerateConfiguration",
                "dynamodb:DeleteTable",
                "s3:PutEncryptionConfiguration",
                "s3:GetBucketObjectLockConfiguration",
                "s3:GetObjectAcl",
                "s3:GetEncryptionConfiguration",
                "dynamodb: TagResource",
                "dynamodb:DescribeTable",
                "s3:GetBucketRequestPayment",
                "dynamodb:GetItem",
                "dynamodb:DescribeContinuousBackups",
                "s3:DeleteBucket",
                "s3:PutBucketVersioning",
                "dynamodb:ConditionCheckItem",
                "dynamodb: UntagResource",
                "s3:GetBucketWebsite",
                "dynamodb:Scan",
                "dynamodb:Query",
                "s3:GetBucketVersioning",
                "dynamodb: DescribeTimeToLive",
                "s3:GetBucketAcl",
                "s3:GetReplicationConfiguration",
                "dynamodb:CreateTable",
                "s3:GetObject",
                "s3:GetBucketCORS",
                "dynamodb:DescribeBackup",
                "s3:GetBucketLocation",
                "s3:GetObjectVersion",
                "dynamodb:GetRecords"
            ],
            "Resource": [
                "arn:aws:dynamodb:MY-REGION:ACCOUNT-ID:table/MY-
DYNAMODB-TABLE",
                "arn:aws:dynamodb:MY-REGION:ACCOUNT-ID:table/MY-
DYNAMODB-TABLE/backup/*",
                "arn:aws:dynamodb:MY-REGION:ACCOUNT-ID:table/MY-
DYNAMODB-TABLE/index/*",
                "arn:aws:dynamodb:MY-REGION:ACCOUNT-ID:table/MY-
DYNAMODB-TABLE/stream/*",
                "arn:aws:s3:::MY-S3-BUCKET",
                "arn:aws:s3:::MY-S3-BUCKET/linuxvm/terraform.tfstate",
                "arn:aws:s3:::MY-S3-BUCKET/ami/terraform.tfstate"
    ]
}
```





Replace ACCOUNT-ID, MY-REGION, MY-DYNAMODB, and MY-S3-BUCKET with their appropriate values for your environment.

- 2. Create an IAM policy for both the Linux VM and AMI stack.
 - a. Navigate to Services, and then IAM, and then Policies, and then Create policy.
 - b. Select the JSON tab.
 - c. Paste the following JSON into the JSON editor.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "VisualEditor0",
            "Effect": "Deny",
            "Action": [
                "ec2:TerminateInstances",
                "ec2:RunInstances"
            ],
            "Resource": "arn:aws:ec2:MY-REGION:ACCOUNT-ID:instance/*",
            "Condition": {
                "StringNotEquals": {
                    "ec2:InstanceType": "t2.micro"
        },
            "Sid": "VisualEditor1",
            "Effect": "Allow",
            "Action": [
                "ec2:DeregisterImage",
                "ec2:DescribeInstances",
                "ec2:ImportKeyPair",
                "ec2:DescribeTags",
                "ec2:DescribeSnapshotAttribute",
                "ec2:DescribeInstanceAttribute",
                "ec2:RegisterImage",
                "ec2:DescribeSnapshots",
                "ec2:DescribeInstanceCreditSpecifications",
                "ec2:DescribeSecurityGroups",
                "ec2:DescribeVolumeAttribute",
                "ec2:DescribeImages",
                "ec2:DescribeVolumeStatus",
                "ec2:DescribeVpcs",
                "ec2:DescribeVolumes",
                "ec2:DescribeAccountAttributes",
               "ec2:DescribeSubnets",
                "ec2:DescribeKeyPairs",
                "ec2:DeleteKevPair",
                "ec2:DescribeInstanceStatus"
            ],
            "Resource": "*"
```



```
},
    "Sid": "VisualEditor2",
    "Effect": "Allow",
    "Action": [
        "ec2:DetachVolume",
        "ec2:AttachVolume",
        "ec2:CreateTags"
    ],
    "Resource": [
        "arn:aws:ec2:*:ACCOUNT-ID:network-interface/*",
        "arn:aws:ec2:*:ACCOUNT-ID:security-group/*",
        "arn:aws:ec2:MY-REGION::snapshot/*",
        "arn:aws:ec2:*:ACCOUNT-ID:subnet/*",
       "arn:aws:ec2:MY-REGION:ACCOUNT-ID:volume/*",
       "arn:aws:ec2:*:ACCOUNT-ID:instance/*",
        "arn:aws:ec2:MY-REGION:ACCOUNT-ID:key-pair/*",
        "arn:aws:ec2:*::image/*"
   1
},
    "Sid": "VisualEditor3",
    "Effect": "Allow",
    "Action": [
        "ec2:TerminateInstances",
        "ec2:RunInstances"
    "Resource": [
        "arn:aws:ec2:MY-REGION:ACCOUNT-ID:subnet/*",
        "arn:aws:ec2:MY-REGION::image/ami-*",
        "arn:aws:ec2:MY-REGION:ACCOUNT-ID:instance/*",
        "arn:aws:ec2:MY-REGION:ACCOUNT-ID:volume/*",
        "arn:aws:ec2:MY-REGION:ACCOUNT-ID:network-interface/*",
        "arn:aws:ec2:MY-REGION:ACCOUNT-ID:key-pair/*",
        "arn:aws:ec2:MY-REGION:ACCOUNT-ID:security-group/*"
},
    "Sid": "VisualEditor4",
    "Effect": "Allow",
    "Action": [
        "s3:PutObject",
        "s3:GetObject",
        "ec2:DeleteVolume",
        "ec2:DeleteSnapshot",
        "dynamodb:PutItem",
        "ec2:CreateKeyPair",
        "dynamodb: DeleteItem",
        "ec2:CreateSnapshot",
        "s3:ListBucket",
        "dynamodb: UpdateItem",
        "ec2:CreateVolume"
    "Resource": [
        "arn:aws:s3:::MY-S3-BUCKET",
        "arn:aws:s3:::MY-S3-BUCKET/*",
```



Replace ACCOUNT-ID, MY-REGION, MY-DYNAMODB, and MY-S3-BUCKET with their appropriate values for your environment.

- 3. Create a t2.micro EC2 Linux instance from which to run terraform. See the AWS documentation for instructions.
- 4. Apply the IAM policies created above to the EC2 instance.
- 5. From your EC2 instance, install the latest version of terraform.
 - a. Navigate to https://releases.hashicorp.com/terraform/
 - b. Download the latest version.
 - c. Unzip the zip file.

```
unzip terraform * linux amd64.zip
```

d. Copy the terraform binary to a location within your environment's PATH variable.

```
sudo cp terraform /usr/local/bin/
```

6. Create an SSH keypair and set the permissions on the private key.

```
ssh-keygen -t rsa -b 4096 -f aws-acme-key chmod 400 aws-acme-key
```

7. Download the SBC image from the Oracle Software Delivery Cloud. The zip file contains the qcow2 image. Download the Stackbuilder files from your SBC version's public Documentation Page. Copy the terraform stack builder file and the qcow2 image to your EC2 instance.

```
scp -i .ssh/aws-key nnSCZ840p3-img-vm_kvm.tgz ec2-user@10.1.1.1:
scp -i .ssh/aws-key nnSCZ840p3_tfStackBuilder.tar.gz ec2-user@10.1.1.1:
```

8. Extract the terraform scripts and the qcow2 image.

```
tar xvf nnSCZ840p3-img-vm_kvm.tgz
tar xvf nnSCZ840p3_tfStackBuilder.tar.gz
```



Create an S3 Bucket

Follow the procedure below if you want to create an S3 bucket. If you have an existing S3 bucket you want to use, skip this section.

1. Navigate to the awk/AMI/s3 directory.

```
cd aws/AMI/s3
```

Initialize terraform.

```
terraform init
```

3. Get terraform packages and validate the templates.

```
terraform get
terraform validate
```

4. Plan and apply to create an S3 bucket DynamoDB table.

```
terraform plan
terraform apply -auto-approve
```



Save the state file to the same bucket for later creating the S3 stack using AWS console

Create the AMI

Three files within the aws/AMI folder contains variables in the standard format of variableName = "value". Each variable within the three variable files should be updated to reflect your environment.

Navigate to the aws/AMI directory.

```
cd ~/aws/AMI
```

- 2. Update the variables in the following files with suitable values:
 - linuxvm/terraform.tfvars
 - linuxvm/bucket.tf
 - ami/bucket.tf
- 3. Run the bash script <code>create_ami.sh</code> with the <code>-c</code> argument to specify the workspace name.

```
bash create ami.sh -c nnSCZ840
```

The script runs a series of terraform commands to create the AMI.

Refer to the README file for more information.

For instructions on using an AMI, see Use an AMI in the AWS documentation.

Delete the AMI

If you no longer need the AMI, follow these steps to delete it.

Navigate to the aws/AMI/ directory.

cd aws/AMI

2. Run the create ami.sh script with the -d flag.

bash create ami.sh -d nnSCZ840

Create and Deploy on OCI

You can deploy the Oracle Communications Session Border Controller (SBC) on Oracle Cloud Infrastructure (OCI) in either standalone or High Availability (HA) mode. When deployed on this platform, you configure and operate the SBC as you would on any other platform. You deploy the SBC to use the environment's IP infrastructure, including the private and public addressing scheme, and its translation functions to protect the OCI management environment from direct exposure to the SBC service delivery environment. These deployments also use OCI's DHCP to establish consistent, compliant management addressing and flat networking across both management and service traffic.

Prerequisites to Deploying an OCI Instance

The OCI deployment infrastructure provides a flexible management system that allows you to create objects required during the instance deployment procedure prior to or during that deployment. When created prior to deployment, these objects become selectable, typically from drop-down lists in the appropriate deployment dialogs. You may use these objects for a single deployment or for multiple deployments.

Deployment prerequisite tasks:

- Identify and deploy to the correct OCI Region. This is typically a default component of your OCI Account.
- Identify and deploy to the correct OCI Availability Domain. By deploying 2 (HA) instances
 during deployment at the same time, you are ensuring that both instances either reside in
 the same Availability Domain or are attached to the same regional subnet if they are
 located in different Availability Domains.
- Identify and deploy to the correct OCI Fault Domains (HA only). You deploy HA instances
 in the same Availability Domain, and in separate Fault Domains.
- Create an Oracle Virtual Cloud Network (VCN). Required VCN configuration includes:
 - Security list—These access control lists provide traffic control at the packet level.
 - Subnet configuration—The SBC has 3 types of vNICs, including management (wancom0), HA (wancom1/wancom2) and Media (s0p0, s1p0 etc). To maintain traffic separation, each of the vNICs should be connected to a separate subnet within the VCN.
 - Internet Gateway—Create a default internet gateway for the compartment and give it an appropriate name.



- Route table (Use Default)—Create a route table to route appropriate Subnet(s) through the Internet Gateway.
- DHCP options (Use Default)—Enable DHCP on the VCN by creating a set of DHCP options, and using the default resolver.

There are additional VCN components that you may find useful for your SBC deployment. These include:

- Dynamic Routing Gateway
- Local Peering Gateways
- NAT Gateways
- Service Gateways

Create Security Lists

Security lists specify the type of traffic allowed on a particular type of subnet. SBC deployments typically need 2 lists, but you may use three if there are specific rules that apply to your HA subnet and are different from your management subnet.

Rules set on security lists can be either stateful or stateless. Stateful rules employ connection tracking and have the benefit of not requiring exit rules. However, there is a limit to the number of connections allowed over stateful connections. and there is a performance hit. Oracle, therefore, recommends stateless lists for media interfaces.



The SBC implements its own ACLs. Protocol access may require that you configure OCI security lists and SBC ACLs. In addition, the port numbers you use within SBC ACLs should match those configured in these security groups.

The security list for management ports can be stateful. Ports you should consider opening for management interfaces include:

- SSH—TCP port 22
- NTP—UDP port 123
- SNMP—UDP port 161
- SNMP Trap—UDP port 162

The security list for media ports should be stateless. Ports you should consider opening for management interfaces include:

- SIP—UDP or TCP port 5060
- SIP TLS—TCP port 5061
- H323—TCP port 1719
- RTP —UDP or TCP port 5004 and 5005

Oracle recommends using a private subnet for HA and a basic security list that allows all local traffic. However, there are some deployments where this is not possible. In these cases, create a security list with a port open for the port you've selected in redundancy-config, which is typically port 9090.



Create Networks and Subnets

OCI interface types include those hidden from the internet and those that are not. In addition, if you are deploying the SBC in HA mode, you must ensure that the cloud can switch between media interfaces on HA instances during failover. This requires secondary private and reserved public addressing. The table below lists configuration requirements and considerations for interfaces, with respect to OCI interface types.

other traffic comes for through Internet. Media through Internet. interfaces) (N/A for HA) (Optional for through standalone (Can be used instead	vNIC Subnet	Public or Private	Required for Standalone	Required for HA	Private IP	Public IP - Ephemeral	Secondar y Private IP	Reserved Public IP
ephemeral	wancom1 wancom2 s0p0, s1p0, s0p1 (and all other Media	Private Private	N/A N/A Between 1 and 8	Required Optional Between 1 and 8	Required Required	N/A N/A Yes for standalone mode, if traffic comes through Internet.	N/A N/A Yes for HA mode. (Optional for	N/A N/A Yes for HA mode, if traffic comes through Internet. (Can be used instead of

Oracle recommends creating regional subnets, which means the subnet can span across availability domains within the region. With this primary and secondary SBC instances can be deployed in two different Availability Domains thereby making use of OCI infrastructure level high availability. Alternatively you could create non-Regional subnets which means the subnet is limited to a single Availability Domain. In this case, both primary and secondary SBC instances MUST be deployed within that Availability Domain

Refer to OCI's Regional Subnets documentation for further information about using these objects.

During the deployment procedure, ensure that OCI provides the IP address for the wancom0 (primary management) interface via DHCP.

Create Dynamic Group and Policy Statements

High Availability (HA) instances require the ability to interact with the platform's API during failover events. You create both Dynamic Groups and Policy Statements for this purpose. Dynamic Groups include rules configuration, which you use to define the instances that belong the group. Policy statements refer to dynamic group names, followed by the action allowed for the group.

The use of dynamic groups allows you to provide the required privileges to multiple instances at the same time. This means you can define and provide these privileges to all instances in a compartment simultaneously by specifying the compartment ID as a single rule attribute. If you want to limit the privileges to a subset of the instances in a compartment, you need to add rule attributes accordingly. If you want to specify individual instances, you need to create the instance first so you have its OCID available as an additional attribute.



Create Dynamic Group

To create a Dynamic Group:

- 1. From the Oracle Cloud **VCN Compartment** dialog, click the Hamburger menu icon to display its drop-down menu and click **Identity**, **Dynamic Groups**.
- Name—Enter a name for your dynamic group. This name can be anything. You use this name when configuring policy statements.
- 3. Matching Rules—Use matching rules configuration to define which instances belong to your group. You specify the attributes on which you want to group instances, and add them to the dynamic group. You can use the provided rule builder, which automatically populates the matching rule text box with your rule(s).
 - Consider the identifier syntax, <instance.compartment.id.> Be sure to use the OCID of the compartment in which SBC you create your instances.



Your IAM policy must include you in the Administrators group to manage dynamic groups.

Create Policy Statements

To create applicable policy statements:

- From the Oracle Cloud VCN Compartment dialog, click the Hamburger menu icon to display its drop-down menu and click Identity, Policies.
- Policy Name—Enter a name for your policy. This name can be anything.
- Policy Statements—Create the statements needed by the instances in your dynamic group to perform HA procedures with the platform's API, including:
 - Allow dynamic-group <dynamic-group name> to read all-resources in compartment
 <compartment name>
 - Allow dynamic-group < dynamic-group name> to use private-ips in compartment <compartment name>
 - Allow dynamic-group <dynamic-group name> to use vnics in compartment <compartment name>
 - Allow dynamic-group < dynamic-group name> to use vnic-attachments in compartment <compartment name>

For <dynamic-group-name>, use the same name you used for the dynamic group you created above.

Deploying the OCI Instance

The OCI instance configuration procedure includes a multi-dialog wizard that presents configuration options in the preferred sequence. The result of this wizard is an installed, operational SBC or two HA SBC instances with management networking only. Having completed the pre-requisites, OCI is able to display objects, including network/subnets and Security Groups, for you to simply select during deployment.

Deploying the SBC on OCI, whether or not you are using HA, includes the following high-level steps:

- Upload Image
- Select Shape
- Create Instance
- 4. Attach the Network Interfaces to Your Instances—Check your interface assignments using the SBC ACLI interface-mapping commands after SBC startup and correct interface assignment, if necessary. See this Oracle® Communications Session Border Controller Platform Preparation and Installation Guide for further instructions on using these commands.
- 5. Create a Console Connection.
- 6. Apply the SBC Configuration—This is an SBC ACLI configuration procedure.

The OCI workspace may present dialogs and fields that differ from this procedure. For full information on deploying OCI instances, see the Oracle OCI documentation.

Upload Image

Upload the SBC disk file to a **Bucket** and create a Custom Image that you can then use when creating an SBC instance. Note that the file you upload must be in **qcow2** format

- Click the Hamburger menu to display the drop-down menu and click Storage > Buckets.
- If there is no Bucket available to store the disk file in the Compartment, create one by clicking the Create Bucket button.
- 3. After creating your **Bucket**, select it and click the **Upload Object** button.
- 4. From the Upload Object dialog, locate your gcow2 image and upload it to the bucket.
- From the overflow menu, which appears as three dots (elipses) at the bottom right corner of the bucket's dialog, click Create a pre-authenticated request for access to the stored object from the Custom Image.
- Enable Read/write permissions by selecting the PERMIT READ ON AND WRITES TO THE OBJECT radio button from the Create Pre-Authenticated Request dialog.
- 7. Click the Create Pre-Authenticated Request button.
- From the Pre-Authenticated Request Details dialog, click the copy link under the PRE-AUTHENTICATED REQUEST URL field and save this URL for accessing the image later.

You use this URL when generating your Custom Image.

Create Image

After uploading the SBC qcow2 disk file and generating a pre-authenticated request, you can import it as a Native OCI image. This process may take 5-10 minutes, after which the Custom Image is available to deploy as an SBC instance.

- From the Oracle Cloud Hamburger menu select Compute, Custom Images.
- From the Import Image dialog, select the correct compartment from the CREATE IN COMPARTMENT drop-down.
- Type a name in the NAME field.
- 4. Select Linux from the **OPERATING SYSTEM** drop-down.
- 5. Select IMPORT FROM AN OBJECT STORAGE URL.
- Paste your PRE-AUTHENTICATED REQUEST URL into the OBJECT STORAGE URL field.



- Select QCOW2 from the IMAGE TYPE radio buttons.
- Select NATIVE MODE from the LAUNCH MODE radio buttons.
- 9. Click Import Image.

If you are deploying a VM.Optimized3.Flex shap, select the created image and perform these last two steps

- In the selected image, select Edit Details and under compatible shapes listed, select VM.Optimized3.Flex
- 2. Select Save change.

Create an SBC Instance

OCI notifies you when it completes the image upload, after which you can create as many SBC instances as desired. You can also choose a marketplace image to create an instance.

- 1. From the Oracle Cloud Hamburger menu, select Compute, Instances .
- 2. On the Instance screen for your Compartment, click the **Create Instance** button.
- 3. From the Create Compute Instance dialog, specify your instance Name.
- 4. Specify the Availability Domain to which you want to deploy your SBC.
- 5. Under the Image, click Change image.
- 6. From the Browse all images dialog, select the Image source as Custom Images.
- 7. From the list, select the desired image.
- Click the Select image button.
- Under Shape, click Change Shape.
- 10. Set Instance type as Virtual Machine.
- 11. Select a supported VM.Standard Shape. After deployment, you can change the shape of your machine by, for example, adding disks and interfaces.
 See your software version's release notes for tables of supported machine shapes.

If you are deploying a VM.Optimized3.Flex shape, select the created image and perform these last four steps

- 1. Set **Shape series** as **Intel**.
- 2. Set Shape name as VM.Optimized3.Flex.
- 3. Modify Number of OCPUs and Amount of memory (GB).
- Click on select shape.

Configure Networking

In the Configure networking section of the Create Compute Instance dialog:

- Select your Virtual Cloud Network Compartment.
- Select a previously created Virtual Cloud Network. This vNIC is the wancom0 interface of the SBC.
- 3. Select a Subnet Compartment.
- 4. Select a Subnet.

The subnet you select for this vNIC should be the one you created for management traffic.

5. Select the **ASSIGN A PUBLIC IP ADDRESS** option.



- Click on Show Advanced Options.
- 7. Under Advanced Options, Launch options, select SRIOV networking.

Add SSH Keys


If needed, you may instantiate your SBC with your own SSH public key configured as an authorized-key for the local admin account.

Copy your client's SSH public key.



On Linux, a user's SSH public key is stored in the <code>.ssh/id_rsa.pub</code> file. If this file does not exist, create it with the command <code>ssh-keygen -t rsa -b 4096</code>.

- 2. In the Add SSH keys section, select Paste SSH Keys.
- Paste your SSH public key into the SSH Keys field.



Boot Volume

If a larger boot volume is desired, select **Specify A Custom Boot Volume Size** and enter the number in gigabytes. You can use the default size or manually specify for your deployment. If manually specifying disk size, use the formula ((2 * RAM) + 12GB) to ensure your disk size is adequate. See the *Platform Preparation and Installation Guide* for extended detail on disk partitions and size.

Creating the Virtual Machine

- 1. Confirm your settings are correct.
- Click Create to instantiate your SBC.
- 3. On the Work Requests page, note the public IP address and the OCID.

Selecting Networks and Subnets

The minimum SBC deployment typically has four interfaces, so at a minimum use **Create vNIC** under Attached VNICs to add:

HA uplink (wancom 1)

- At least 2 Media Interfaces
- 1. Navigate to Compute, and then Instances and open your instance.
- 2. Click the **Start** button to start your instance, if it is not started. The infrastructure cannot add interfaces to an instance when it is in a stopped state.
- Scroll down to the Attached VNICs section of your instance dialog and click the Create VNIC button.
- Within the Create VNIC dialog, name your VNICs and select the subnets you created for them.



The vSBC does not support VNIC deletion on OCI VM.Optimized3.Flex shapes. So interfaces once added cannot be deleted

Create a Console

If needed, create a console for your instance with appropriate SSH keys.

- 1. Scroll to the Console Connections section of the instance, available under Resources.
- 2. Click the Create Console Connection button.
- 3. From the Create Console Connection, either choose or paste in your SSH Key files.
- 4. Click the Create Console Connection

After creating the console connection, you can access it via SSH or VNC.

Set the User and Administrative Passwords on the SBC

These password procedures are required before any further SBC operations. For HA deployments, perform these procedures on both SBCs.

- 1. Click Compute, and then Instances, and open your newly created SBC.
- 2. Under the Instance Iniformation tab, copy the OCID.
- When the virtual machine has finished initializing, SSH to the public hostname. The username is "user" and the initial SSH password is "acme" + the OCID.
- Set the user password by logging in for the first time.



5. Set the administrative password by typing **enable** at the command prompt. The initial enable password is "packet" + the OCID.

```
> enable
Password:
*ALERT*
*****************
admin password has not been set. Please set password now.
******************
** Only alphabetic (upper or lower case), numeric and punctuation
** characters are allowed in the password.
** Password must be 8 - 64 characters,
** and have 3 of the 4 following character classes:
** - lower case alpha
** - upper case alpha
** - numerals
** - punctuation
****************
Enter New Password:
Confirm New Password:
```

Verify the network interfaces have MAC addresses.
 Use the show interfaces mapping command to verify the network interfaces have MAC addresses.

The **interface-mapping** branch includes the **swap** command, which allows you to correct interface to MAC address mappings.

```
Interface Mapping Info
Eth-IF MAC-Addr
                               Label
wancom0 06:DF:71:BA:D8:77  #generic
wancom1 06:A6:08:58:92:C9  #generic
s0p0 06:D4:E6:E8:B8:FB  #generic
s1p0
      06:EA:08:51:4D:DF
                                #generic
wancom2 FF:FF:FF:FF:FF
                               #dummy
spare FF:FF:FF:FF:FF
                                #dummy
s0p1
        FF:FF:FF:FF:FF
                                 #dummy
s1p1 FF:FF:FF:FF:FF
                                 #dummy
s0p2 FF:FF:FF:FF:FF
                                 #dummy
s1p2 FF:FF:FF:FF:FF
                                 #dummy
s0p3
      FF:FF:FF:FF:FF
                                 #dummy
s1p3 FF:FF:FF:FF:FF
                                 #dummy
```

7. Reboot the virtual machine.

show interfaces mapping

reboot



Refer to the *Oracle® Communications Session Border Controller Configuration Guide* after you have completed this deployment for administrative and service configuration, including product setup, entitlement setup and HA configuration.

Migrating to an Optimized3. Flex Machine Shape


The SBC supports machine migration to and from supported versions on OCI VM.Optimized3.Flex shapes.

The table lists the supported standard shapes that you can migrate to their respective VM.Optimized3.Flex shapes. Changing instance type from VM.Standard2.X to VM.Optimized3.Flex is supported only when both the instances have the same number of OCPUs, the same amount of memory, the same number of vNICs attached, and the same type of networking mode (SRIOV).


From Instance Type	To Instance Type	OCPUs	Memory
VM.Standard2.4	VM.Optimized3.Flex- Small	4	60
VM.Standard2.8	VM.Optimized3.Flex- Medium	8	120
VM.Standard2.16	VM.Optimized3.Flex- Large	16	240

The SBC supports this instance type modification for both standalone and HA deployments. Follow the steps below to migrate a standalone VM from the VM.Standard2.X shape to the VM.Optimized3.Flex shape .


- Upgrade the software on VM.Standard2.X shape to the vSBC version that supports VM.Optimized3.Flex shapes.
- After upgrading the software, stop your VM from the OCI console.
 The image below depicts a VM.Standard2.4 instance's Shape Configuration before migration.




Stop your instance by clicking the "stop" button.



Navigate to Edit, Edit Shape, and select the target shape name (Intel, VM.Optimized3.Flex).



Set the same number of OCPUs and amount of memory as that of standard shape, then Save changes.



5. Start your migrated vSBC machine.

Migrating an HA Pair

Follow the steps below to migrate an HA pair of VM.Standard2.X instances to VM.Optimized3.Flex shapes. Both the starting and target instances of the HA pair must be the same instance type, with the same number of OCPUs, the same amount of memory, and the same networking type (SRIOV).

- 1. Upgrade the software on both nodes of VM.Standard2.X shapes to the vSBC version that supports VM.Optimized3.Flex shapes.
- 2. From OCI console, stop your standby VSBC.



- Navigate to Edit, Edit Shape, and select the target shape name (VM.Optimized3.Flex).
- 4. Set number of OCPUs and amount of memory to be the same as that of standard shape and save your changes.
- Start your migrated standby vSBC machine.
- **6.** After startup, check that the original active and standby are in the same roles and are synchronized.
- From OCI console, stop your active vSBC. The vSBC over the new Flex VM takes the role as new Active.
- 8. Navigate to Edit, Edit Shape, and select the target shape name (VM.Optimized3.Flex).
- Set the number of OCPUs, the amount of memory (the same as that of standard shape) and save your changes
- **10.** Start the instance. This instance takes the new standby role.
- 11. Verify the HA roles and that the systems are synchronized.

If you want to change the number of forwarding cores in your deployed VM.Optimized3.Flex shape, change the active instance first and reboot it after the configuration fully synchronizes. Once the first VM comes up, it takes the role as "Standby". Reboot the newly active vSBC for the forwarding core count changes to come into effect on both instances of your HA pair.

Create and Deploy on OCI using Resource Manager

OCI Resource Manager automates the process of provisioning your Oracle Cloud Infrastructure resources. The Resource Manager provides stacks to set up OCI resources that runs the virtual SBC using Terraform scripts. However, Terraform scripts cannot be used for complete SBC configuration. Hence, Resource Manager uses two pre-build stacks for deploying environments. The two stacks are - VCN and SBC stack. The VCN stack creates the required network infrastructure to deploy the virtual SBC instance on OCI. The SBC stack instantiates a standalone or HA pair on OCI with all Day-0 configuration, for example: loading product type and entitlements, configuring cores, setting up HA configuration, SNMP etc.

Workflows that can be automated with Stacks include creating:

- Virtual Cloud Network
- Internet Gateway
- Route Tables
- Security Lists or Network Security Groups for Management Interface, Media Interface and HA
- Subnets
- VNIC and assigning to SBC instance
- Setting up interface mapping
- Setting SNMP and NTP configuration
- Setting HA configuration

Prerequisites for creating VCN and SBC stacks

The minimum supported version of the Terraform is 0.13.



Prerequisites

Prerequisites for creating VCN and SBC stacks

- The minimum supported version of the Terraform is 0.12.
- The VCN and SBC folder is a part of the Terraform stack builder tar file (For example -SCZ910_tfStackBuilder.tar.gz). Download the Terraform stack builder file and extract to obtain the OCI folder. The OCI folder contains VCN and SBC folders.
- Refer the below common files in VCN and SBC folder for more information.
 - README.txt Contains details about prerequisites, installation, usage of variable and running the template.
 - variable.tf file Contains variable description and their default value.
 - versions.tf file Contains details about the minimal version supported by the Terraform,
 OCI provider and any other providers used by the templates.



It is recommended to avoid special characters while editing the entitlement parameters.

- Choose a unique Resource Label for a deploying a new SBC or VCN stack.
- Choose an instance shape based on the number of cores, RAM size and vNICs to attach to a SBC stack.
- The Password for user and admin account must be SHA1 encoded while configuring a SBC.

Prerequisites for creating SBC stack

To deploy a standalone SBC or HA pair on OCI, create the below list of network resources using a VCN stack.

OCI Resources	Quantity	Notes
Virtual Cloud Network	1	NA
Internet Gateway	1	Create an internet gateway only if management or media subnets are chosen as public subnets.
Service Gateway	1	Create only if wancom0 subnet is defined as private.
Subnets	7	 1 – wancom0 (Private or Public)
		 2 – Private subnets (wancom1 & wancom2)
		 4 – Media Subnets for media S0P0, S1P0, S0P1 S1P1 (Private or Public)



OCI Resources	Quantity	Notes
Route Tables	3	Management RTHA RTMedia RT
Security Lists	3	Management Security ListHA Security ListMedia Security List


Creating VCN Stack

A VCN stack creates the required network infrastructure to deploy the virtual SBC instance on OCI.

To create a VCN stack using OCI Resource Manager:

- From the Navigation menu, click Oracle Cloud Developer Services.
- 2. Under Resource Manager, click Stacks.
- 3. In the Create Stack dialog, under Choose the origin of Terraform configuration, select the option My Configuration.
- 4. Under Stack Configuration, select Folder option,
 - Drag and drop a folder or click Browse and navigate to the location of the VCN folder.
- 5. Upload the VCN folder.
- 6. Enter a **Name** for the stack or accept the default name.
- 7. Enter a Description.
- 8. Select the **Compartment** you want to create the stack.
- Apply tags to the stacks, if required.
- 10. Click Next.






- **11.** From the **Configure Variables** panel, select the variables from the Terraform configuration.
 - a. Under VCN Configuration, enter the Resource Label.

Note:
The Resource Label precedes every resource name in the VCN.

- **b.** Enter the **VCN CIDR** or accept the default value provided.
- Under Object Storage Bucket, select Create OCI storage Bucket to create a bucket or container to deploy SBC.
- 13. Enter a Bucket Name or accept the default provided.




14. Under Management Interface Configurations,

- Deselect Prohibit Public IP on Wancom0 subnet to create public subnet for management interface.
- b. Select **Enable SSH** to add ingress rule to allow SSH for management interface.
- c. Select **Enable FTP** to add ingress rule to allow FTP for management interface.
- d. Select **Enable ICMP** to add ingress rule to allow ICMP for management interface.
- e. Select **Enable HTTPS/HTTPS** to add ingress rule to allow HTTP/HTTPS protocol for management interface.

15. Under HA Interface Configurations,

- a. Enter the IP address range for Wancom1.
- b. Enter the IP address range for Wancom2.






16. Under Media Interface Configurations,

- a. Enter the IP address range for S0P0 media interface.
- b. Enter the IP address range for S1P0 media interface.
- c. Enter the IP address range for S0P1 media interface.
- d. Enter the IP address range for S1P1 media interface.
- e. Select **Enable ICMP** to add ingress rule to allow ICMP for media interface.
- f. Select **Enable TCP** to add ingress rule to allow TCP for media interface.
 - For TCP Port Min, enter the minimum TCP port value.
 - For TCP Port Max, enter the maximum TCP port value.
- g. Select **Enable UDP** to add ingress rule to allow UDP for media interface.
 - For UDP Port Min, enter the minimum UDP port value.
 - For UDP Port Max, enter the maximum UDP port value.






17. Click Next.

From the **Review** panel, verify the configuration variables.




You can only view the variables whose default values have changed.





- 18. Click Create to create the VCN stack. The new stack details appear. If you click Plan, you can view the resources that will be created based on the configuration variables.
- 19. Click Apply to create the OCI resources.



Once the VCN stack is created, you can view the VCN from the Navigation Menu > Networking > Virtual Cloud Networks > Virtual Cloud Networks in <your> Compartment.

Creating SBC Stack

A SBC stack instantiates a standalone or HA pair on OCI with Day-0 configurations, for example: loading product type and entitlements, configuring cores, setting up HA configuration, SNMP etc.

The Day-0 configurations include:


- Configuring product type and entitlements.
- Configuring default REST interface (using self-signed certificates).
- Configuring network mapping.
- Configuring disk partitions.
- · Changing user/admin password.
- Configuring CPU core assignments.
- Configuring SNMP/NTP service.
- Redundancy configuration for HA.

To create a SBC stack:

- From the Navigation menu, click Oracle Cloud Developer Services.
- 2. Under Resource Manager, click Stacks.
- 3. In the Create Stack dialog, under Choose the origin of Terraform configuration, select the option My Configuration.
- 4. Under Stack Configuration, select Folder option,



- Drag and drop a folder or click Browse and navigate to the location of the SBC folder.
- 5. Upload the SBC folder.
- 6. Enter a Name for the stack or accept the default name.
- 7. Enter a **Description**.
- 8. Select the **Compartment** you want to create the stack.
- 9. Select the Terraform version.
- **10.** Apply **tags** to the stack, if required.
- 11. Click Next.



- **12.** From the **Configure Variables** panel, select the variables from the Terraform configuration.
 - a. Under SBC General Configurations, enter the Resource Label.






The **Resource Label** precedes every resource name in the SBC.

- b. Enter the Bucket Name.
- c. Select the Product Type.
- d. Select Enable High Availability to create HA pair.
- e. Optionally, select Enable Rest Interface to enable the SBC ports to serve REST API requests with self-signed certificates.
- f. Select Apply SBC Base Configuration to apply Day-0 configurations while creating the SBC.
- g. Select Create Data Disk Partitions to create default disk partitions.



The default disk partitions are: /app - 20% of memory and /sys - 80% of memory.

- h. Select the Licenses .
- i. Optionally enter the User Password.
- j. Optionally enter the Admin Password.






13. Under Instance Configurations,

- a. Optionally select the SBC Image from the list.
- **b.** Select **Enter SBC Image OCID Manually**, if you want to enter an OCID for SBC Image manually.
- c. Select an Instance Shape.



Instance Shape determines the number of CPUs, amount of memory, and other resources allocated to a new instance. SBC only supports VM Standard 2.x shapes.


- d. Optionally, enter a SSH Public Key
- e. Optionally, enter a Boot Volume Size .
- f. Optionally, enter a **Availability Domain** that is available based on the OCI region.
- g. Optionally, select a Fault Domain, if you choose to create a standalone SBC.



When deploying a vSBC on the VM.Optimized3.Flex shapes, be sure to configure the following fields to these values, as displayed in the dialog below:

- a. Set the Instance Shape to VM.Optimized3.Flex.
- **b.** Set the **Number of OCPUs** to 4.
- c. Set the Amount of Memory (GB) to 16.





14. Under vNIC Configurations,

- a. Select the Virtual Cloud Network from the list.
- b. Select the Wancom0 subnet OCID.
- c. Optionally, enter a static IP address for wancom0 interface (standalone or primary node) in theWancom0 Primary Node IP.
- **d.** Optionally, enter a static IP address for wancom0 interface (secondary node) in the **Wancom0 Secondary Node IP**.
- e. Select Attach HA interface (wancom1) to add a HA interface.
 - Select the Wancom1 subnet OCID.
 - Optionally, enter a static IP address for wancom1 interface (standalone or primary node) in Wancom1 Primary Node IP.
 - Optionally, enter a static IP address for wancom1 interface (secondary node) in Wancom1 Secondary Node IP.



Similarly for wancom2, add the HA interface.

- Select the Attach Media Interfaces (S0p0) to add a media vNIC.
 - Select the S0P0 Subnet OCID.
 - Select Assign Public IP to assign a public IP address to the media interface. This
 must a public subnet.
 - Optionally, enter a static IP address for S0P0 interface (standalone or primary node) in **S0P0 Primary Node IP**.




- Optionally, enter a static IP address for S0P0 interface (secondary node) in **S0P0 Secondary Node IP**.
- Optionally, select a value for static IP address in the S0P0 Static Virtual IPs
- Optionally, enter the S0P0 Count of Virtual IPs for the media interface.



Similarly for other media interfaces, select the static IP address for primary and secondary nodes.





15. Under CPU Core Configurations, optionally


- a. Select the number of CPU Forwarding Core dedicated for forwarding frames.
- b. Enter the number of CPU **DOS Cores** dedicated for Denial-of-Service Protection.
- c. Enter the number of CPU Transcoding Cores dedicated for transcoding media.
- d. Select Use Sibling Cores, if you want to use sibling core for data-path and virtual SBC uses SMT topology.





The sum of number of Forwarding Core, DOS Cores, Transcoding Core and two Signaling Cores must be less than or equal to the number of virtual CPUs.

- 16. Under SNMP Configurations, optionally
 - a. Enter the SNMP Community Name used for SNMP management.
 - b. Enter the SNMP IP Address.
- 17. Under NTP Configurations, optionally
 - a. Enter the **NTP Server 1** used for time synchronization.
 - b. Enter the NTP Server 2.




18. Click Next.

From the **Review** panel, verify the configuration variables.



You can only view the variables whose default values have changed.

19. Select **Run Apply** to automatically provision the OCI resources, when the SBC stack is created.




- 20. Click Create to create the SBC stack.
 The new stack details appear. If you click Plan, you can view the resources that will be created based on the configuration variables.
- 21. Click **Apply** to create the OCI resources.

Destroying Resources

You can destroy resources from Resource Manager.

Click **Destroy** from the **Resource Manager > Stacks > Stack Details**. In case, a SBC is deployed using a VCN, you cannot destroy the VCN.



Running templates from CLI

You can run templates from CLI, similar to running the Terraform templates from OCI Resource Manager.

To run the templates from CLI:

- Download Terraform from the following URL and install Terraform. https://www.terraform.io/downloads.html
- Download the Terraform stack builder file.



The VCN and SBC folder is a part of the Terraform stack builder tar file (For example - SCZ910_tfStackBuilder.tar.gz).

Extract the Terraform stack builder file.

```
tar -xzvf <filename> tfStackBuilder.tar.gz
```

4. Change directory to point to the SBC folder.

cd oci/SBC



If you are creating a VCN, change directory to VCN folder.

5. Uncomment these lines in the provider.tf file.

```
user_ocid = var.user_ocid
fingerprint = var.fingerprint
private_key_path = var.oci_api_private_key_path
```

- 6. Save the file.
- 7. Export the required environment variables. For example:

```
export TF_VAR_region="us-phoenix-1"
export
TF_VAR_tenancy_ocid="ocid1.tenancy.oc1..aaaaaaaauhdcynzumstqiwxbx5d2mkmeqfly
fnwsta6gn4g7ofmnfkq"
export
TF_VAR_compartment_ocid="ocid1.compartment.oc1..aaaaaaaaa6pylpay5csm5xw2e5tz
ine24vtvqpguepm7g45r4rrr7cla"
export TF_VAR_oci_api_private_key_path="$HOME/.oci/oci_api_key.pem"
export
TF_VAR_user_ocid="ocid1.user.oc1..aaaaaaa7f76hacxruga2hjsgp2vqsypx7ejaqrbv54
vgxhymnq"
export TF_VAR_fingerprint="cb:b6:86:6:6d:9c:8a:88:3b:85:72:f9:47:34:5c:43"
```

- 8. Edit the **terraform.tfvars** file and change the configuration variables as required.
- 9. Save the file.
- Edit the entitlements/product>.yaml file and change the entitlement parameters as required.
- 11. Save the file.



12. Run this command to initiate Terraform.

```
terraform init
```

13. Run this command to verify the templates and new OCI resource that will be created based on the configuration variables.

```
terraform plan
```

14. Run this command to create the resources.

```
terraform apply
```

15. Run this command to destroy the resources.

```
terraform destroy
```

16. Change directory to point to the VCN folder.

```
cd oci/VCN
```

Repeat steps 5 through 15 to create a VCN stack.

IAM Policies

To use Resource Manager from OCI console you need to have the following IAM policies.

```
Allow group <group-name> to manageorm-family in compartment <compartment-name>
```

IAM permissions to run VCN stack

You need to have the following permissions to run VCN Stack.

At a broader level, you need these permissions.

```
Allow group <group-name> to manageobjects-family in compartment <compartment-name>
Allow group <group-name> to managevirtual-network-family in compartment <compartment-name>
```

At a granular level, you need these permissions.

```
Allow group <group-name> to manage objectstorage-namespaces in compartment<compartment-name> Allow group <group-name> to manage buckets in compartment<compartment-name> Allow group <group-name> to manage vcns in compartment<compartment-name> Allow group <group-name> to manage subnets in compartment<compartment-name> Allow group <group-name> to manage route-tables in compartment<compartment-name> Allow group <group-name> to manage network-security-groups in compartment<compartment-name> Allow group <group-name> to manage security-lists in compartment<compartment-name> Allow group <group-name> to manage dhcp-options in compartment<compartment-
```

```
name>
Allow group <group-name> to manage internet-gateways in
compartment<compartment-name>
Allow group <group-name> to manage nat-gateways in compartment<compartment-
name>
Allow group <group-name> to manage service-gateways in
compartment<compartment-name>
Allow group <group-name> to manage local-peering-gateways in
compartment<compartment-name>
Allow group <group-name> to manage drgs in compartment<compartment-name>
Allow group <group-name> to manage private-ips in compartment<compartment-
name>
Allow group <group-name> to manage volume-attachments in
compartment<compartment-name>
Allow group <group-name> to manage instance-console-connection in
compartment<compartment-name>
```

IAM permissions to run SBC stack

At a broader level, you need these permissions.

```
Allow group <group-name> to manage objects-family in compartment<compartment-name>
Allow group <group-name> to manage instance-family in compartment<compartment-name>
```

At a granular level, you need these permissions.

```
Allow group <group-name> to manage objectstorage-namespaces in
compartment<compartment-name>
Allow group <group-name> to manage buckets in compartment<compartment-name>
Allow group <group-name> to manage objects in compartment<compartment-name>
Allow group <group-name> to use private-ips in compartment<compartment-name>
Allow group <group-name> to use public-ips in compartment<compartment-name>
Allow group <group-name> to use vnics in compartment<compartment-name>
Allow group <group-name> to use vnic-attachments in compartment<compartment-
name>
Allow group <group-name> to use subnets in compartment<compartment-name>
Allow group <group-name> to read vcn in compartment<compartment-name>
Allow group <group-name> to read instance-images in compartment<compartment-
name>
Allow group <group-name> to use network-security-groups in
compartment<compartment-name>
Allow group <group-name> to read app-catalog-listing in
compartment<compartment-name>
```

Troubleshooting

The following list describes the troubleshooting steps while creating or destroying a stack.

- While destroying a SBC stack, do not keep the SBC VMs in a stopped state. The SBC becomes unrecoverable and you cannot clean up OCI resources.
- While applying a stack, if a Terraform template throws an error, destroy the stack and delete all resources before re-applying the stack.

- While deleting a stack, if a Terraform template throws an error, re-run the destroy stack action, to completely delete all the resources.
- Navigate to Resource Manager > Stacks > Stack Details > Job Details > Logs to verify an error detail.

Create and Deploy on Azure

You can deploy the Oracle Communications Session Border Controller (SBC) on Azure public clouds. Azure provides multiple ways of managing your environment(s), including via its web portal, using its powershell and its CLI interfaces. This document focuses on the portal. The portal provides navigation via a web-page pane with links to specified functions on the left side of portal pages. These procedure also assume you have reviewed Azure documentation, and can access portal pages and navigation.

Deployment detail specific to Azure includes:

- VM Interfaces—An initial instance deployment creates only a single (wancom0) network interface with an ACL for inbound SSH and HTTP/HTTPS access only. You add additional media interfaces after initial deployment.
- Interface Types—Use Para-Virtual interfaces for both management and media/signaling ports.
- Virtual Hard Disk (VHD) Size—The SBC default disk size for Azure is 20G. Use the formula, (2xRAM)+2+2+8 GB, and reconfigure the disk size during deployment if you plan on using the /sys and /app partitions or if you have larger RAM and disk requirements for your deployment.
- Console Selection—Use the serial console as the default device because Azure does not support a virtualized VGA console at startup.
- Cloud Detection—The SBC show platform command includes the cloud field, which, when applicable, shows the detected cloud environment, including Azure.



Azure auto-fills the target **Subscription**, **Resource Group**, **Region** and **Location** fields. These components are deployment-specific and used in some Azure configuration components. You may need to specify them if your cloud environment includes multiples of these components, which you would have defined prior to instance deployment.

Prerequisites to Deploying an Azure Instance

You can create some of the objects required during the instance deployment procedure prior to or during that deployment. When created prior to instance deployment, these objects become selectable, typically from drop-down lists in the appropriate deployment dialogs. You may use these objects for a single deployment or for multiple deployments.

Tasks denoted here as SBC instance deployment prerequisites include:

- You have identified and are deploying to or via the correct Azure:
 - Subscription
 - Region



- Location
- You have deployed the SBC VHD.
- You have created a bootable image using the Create image dialog.
- You have created the virtual networks you need for media and management interfaces, with the exception of wancom0. Within the context of interface creation, you also need to create:
 - Subnets—The SBC has 3 types of vNICs, including management (wancom0), and Media (s0p0, s1p0 etc). To maintain traffic separation, each of the vNICs should be connected to a separate subnet.
 - Security Groups—These are the equivalent of ACLs, defining the inbound and outbound traffic allowed through that interface.

During the instance interface creation procedure, you must have the appropriate image, subnets and security groups available. You create and attach the wancom0 network (subnet) during instance deployment.

Deploying the SBC VHD File

You acquire the SBC VHD file via your Oracle Support account. You deploy the SBC using the VHD supplied by Oracle, based on the SBC version.

- 1. Create a Blob Container under your Storage Accounts .
- 2. Upload the VHD using the Microsoft Azure Storage Explorer:
 - To a folder under the Blob Container
 - In the same Region where you plan to create your instance.

After the VHD is uploaded, its URI (URL) is available from the Azure GUI. You use this URI as the import source to create a new disk image. You can create instances manually using the Azure GUI, or via script using the PowerShell command-line. This document presents the GUI procedure.

Create Image

After uploading the file, you create an bootable image from the Create image and specify:

- An Image Name.
- An OS disk:
 - Set the OS disk type to Linux.
 - Paste or select your VHD file URI as the Storage blob.
 - Set the Account Type to Standard HDD.
 - Set Host caching to Read/write.

Click the **Create** button at the bottom of the dialog. Azure creates your boot image. You can review this image from your Azure **Home** page.

See Azure's *Upload a generalized VHD and use it to create new VMs in Azure* article for extended instructions on creating your image.

Create Virtual Networks

Virtual networks define the Azure cloud address space within which your interfaces reside. Virtual Subnets divide the virtual networks into segregated subsets in the same way as subnets



that are not virtual. You can assign the actual interface address manually or use Azure's DHCP to acquire it.

You select subnets when you create an instances virtual interfaces. You also select security groups during interface creation.

To create a virtual network:

- From Azure's navigation list on the left side of the portal, click Create a resource, Networking, Virtual network.
- Enter your virtual network's name in the Name field.
- Click Address space and configure the space on the subsequent dialog.
- Repeat these steps for all distinct networks you need for your instance. interfaces.

Creating a virtual network establishes a single virtual subnet that consists of the entire network's address space. You can divide the network into multiple subnets:

- From Azure's navigation list on the left side of the portal, click Virtual networks to display
 a list of all your
- 2. Click the name of your virtual network to access that network's configuration dialog.
- 3. Click Subnets (in the SETTINGS section) on the Create virtual network dialog.
- 4. Click +Add on the subnet's dialog.
- 5. Enter a Name on the **Add subnet** dialog.
- 6. Configure the range in the Subnet address range field.
- Repeat these steps for all the interfaces you need for your instance.

If configured correctly, you can review your subnet configuration on **Subnets** dialog.

See Azure's on line documentation for articles that provide for extended instructions on creating your network elements.

Deploying the Azure Instance

This is the main instance configuration procedure. It includes a multi-dialog wizard that presents configuration options in the preferred sequence. The result of this wizard is an installed, operational SBC with a management interface. You add media interfaces after deployment.

Important requirements include the following:

- Management subnets should be public to allow access from outside the cloud. Addressing
 within these subnets should include Public IP addresses.
- All Media interfaces addresses that must be reachable through the internet must reside on public subnets; all others can reside on private subnets.

The instance deployment wizard sequence includes:

- Basics
- 2. Disks
- Networking
- 4. Managment
- Guest Config



Tags

7. Review and Create

Your Azure workspace may present dialogs and fields that differ from this procedure. For full information on deploying Azure instances, see the Azure documentation.

Basics

During deployment you choose a size for the SBC, based on pre-packaged Azure sizes. After deployment, you can change the detail of these sizes to, for example, add disks or interfaces. Azure presents multiple size options for multiple size types. These size types define architectural differences and cannot be changed after deployment. Azure size types include:

- F(x)—Does not support premium storage
- FS(x)—Supports premium storage
- FS(x)_v2—Supports premium storage and hyperthreading.


See your software version's release notes for tables of supported machine sizing.

On the Basics tab:

- Enter a name as your Virtual machine name.
- Select your Availability Option.
- Specify the Image to deploy.
- Size the image.
- Specifiy your Administrator Account information. The SBC image does not use this
 account, but you must set these fields or the platform does not allow the workflow to
 complete. You may either:
 - Use the authentication type of Password. Create a null user and give it any password. During instantiation, the SBC ignores this password.
 - Use the authentication type of SSH public key. Create a null user and paste in the SSH public key from the machine you'll use to log into the SBC. During instantiation, the SBC will load this key into the ssh-key configuration element as an authorized-key for the admin user.
- Specify your Inbound port rules external port access to interfaces.

The image below displays the initial instance deployment wizard dialog.







Disks

Azure instance deployment **Disk** configuration includes setting the **OS disk type** to **Standard HDD**.

Despite the initial boot disk size provided by Oracle, you are free to create a boot disk that is a different size, as supported by Azure. Specifically, consider whether you need a disk that is larger than the initial size to allow for storing log files and other data. Oracle only supports one primary disk for the SBC, although that disk may have multiple partitions.

Networking

If you plan on using accelerated networking for higher performance, use the Azure CLI to create the network. See Create a network interface with accelerated networking. Even if you create the media ports during the multi-dialog wizard with accelerated networking disabled, you can still enable accelerated networking later in the Azure CLI.

Otherwise, configure fields in the **Networking** dialog to establish the network used by wancom0. Azure instance deployment **Networking** configuration includes:

- Select or create a new Virtual Network.
- Select or create your Subnet.
- Enter a name as your Public IP.
- Select your NIC network security group.
- Select your Select inbound ports.

Leave the **Place this virtual machine behind an existing load balancing solution?** field set to **No**.

Management

Azure instance deployment Management configuration includes:

- Set Boot diagnostics to On.
- Set OS guest diagnotics to On.
- Set Diagnotics storage account to your account.

Leave all other fields set to Off.

Advanced

You do not need to configure anything on the **Advanced** dialog.

Tags

You are free to define tags in any way you want, to help clarify detail about those objects. You do not need to configure anything on the **Tags** dialog.

Review and Create

Use the **Review and Create** to review your settings, then click the **Create** button to complete instance creation.



Create Networking for Additional Interfaces

SBC instance includes establishing networking to the primary management interface, wancom0. You create networking for all other interfaces after deployment. Azure requires that you stop an SBC instance before you create and attach additional network interfaces. If not, Azure displays an error during interface configuration.

At a minimum, create the SBC s0p0 and s1p0 interfaces. This results in having 3 interfaces for the instance, including wancom0. If you create only 3 interfaces for the instance, the default setup will be wancom0, wancom1 and s0p0 at first bootup. Use the ACLI command interface-mapping to remap the interfaces. If you need accelerated networking for higher performance on media interfaces, use the Azure CLI to create the network. See Create a network interface with accelerated networking. Note these two caveats:

- Accelerated networking is not supported for management interfaces.
- If enabled, accelerated networking must be enabled for all media interfaces.

Otherwise, when you select an instance from the portal, Azure displays an instance-specific navigation pane on the left side of the dialog. Begin each interface configuration as follows:

- Click the Settings, Networking link.
- At the top of the interface configuration dialog, click the Attach Network Interface link.
- Within the Attach Network Interface dialog, click the Create Network interface link.

Create Network interface

Configure the applicable **Create Network interface** fields, including:

- Enter a distinguishable name in theName field.
- Select the correct subnet from the Subnet dropdown field.
- Set the Private IP assignment setting to Static.
- Set the **Private IP address** to an address within the subnet.

Security Groups

The **Create Network interface** dialog includes the **Network security group** selection tool. By default, all network interfaces are set to deny all traffic. You assign security groups to each of the media interfaces to specify the traffic you want to allow. Assuming you have created these groups as a pre-requisite to instance deployment, you simply select the appropriate group from **Create Network interface** dialog.

Bear in mind that the SBC has its own traffic control configurations. Ensure that the Azure network interface and SBC configurations do not conflict.

Attach Network Interfaces to the SBC

Upon completing the configuration of the **Create Network interface** fields, you also attach them to your instance. Azure creates MAC addresses upon interface creation. Note these addresses so you can later verify they are attached to the correct SBC, as described below.

Restart the SBC

Start your instance after creating and attaching all interfaces. Use the instance's **Serial Console** to connect to the virtual COM1 serial port. After bootup, proceed with setting your SBC passwords from the SBC command line interface.



Set the User and Administrative Passwords on the OCSBC

This procedure turns to the SBC and performs the password procedures, which are required before any further SBC operations. For HA deployments, perform these procedures on both SBCs.

- Under Azure services, click Virtual machines to locate the newly created SBC.
- Connect to the Cloud Shell.
- 3. Use the az command to retrieve the vmId.

```
az vm show --name <vm-name> --resource-group <resource-group-name>
```

- Note the public IP address.
- 5. When the virtual machine has finished initializing, SSH to the public IP address. The username is "user" and the initial SSH password is "acme" + the vmId.
- 6. Set the user password by logging in for the first time.

```
$ ssh user@<IP address>
user@somewhere.compute-1.azure.com's password:
******************
user password has not been set. Please set password now.
******************
** Only alphabetic (upper or lower case), numeric and punctuation
** characters are allowed in the password.
** Password must be 8 - 64 characters,
** and have 3 of the 4 following character classes:
** - lower case alpha
** - upper case alpha
** - numerals
** - punctuation
******************
Enter New Password:
Confirm New Password:
```

7. Set the administrative password by typing **enable** at the command prompt. The initial enable password is "packet" + the vmId.



```
** - upper case alpha
** - numerals
** - punctuation
Enter New Password:
Confirm New Password:
```

Verify the network interfaces have MAC addresses. Use the **show interfaces mapping** command to verify the network interfaces have MAC addresses.

```
Interface Mapping Info
Eth-IF MAC-Addr
                           Label
wancom0 06:DF:71:BA:D8:77
                          #generic
wancom1 06:A6:08:58:92:C9
                          #generic
s0p0 06:D4:E6:E8:B8:FB
                          #generic
s1p0 06:EA:08:51:4D:DF
                            #generic
wancom2 FF:FF:FF:FF:FF
                            #dummy
spare FF:FF:FF:FF:FF
                            #dummy
s0p1 FF:FF:FF:FF:FF
                            #dummy
s1p1 FF:FF:FF:FF:FF
                            #dummy
s0p2 FF:FF:FF:FF:FF
                            #dummy
s1p2 FF:FF:FF:FF:FF
                            #dummy
s0p3 FF:FF:FF:FF:FF
                            #dummy
      FF:FF:FF:FF:FF
```

show interfaces mapping

The interface-mapping branch includes the swap command, which allows you to correct interface to MAC address mappings.

#dummy

Reboot the virtual machine.

```
# reboot
```

s1p3

Refer to the Oracle® Communications Session Border Controller Configuration Guide after you have completed this deployment for administrative and service configuration, including product setup, entitlement setup and HA configuration.

Deploying on Azure with accelerated networking

Accelerated networking is required for higher performance of media interfaces, but it is not supported by management interfaces. When you create a virtual machine on Azure and login to view the MAC address, the management and media interfaces have incorrect MAC addresses. It is important to swap the interfaces to map them to their proper MAC addresses and enable accelerated networking for all media interfaces.

To create a virtual Oracle Communications Session Border Controller (OCSBC) on Azure with Accelerated networking enabled:

- Create a virtual machine on Azure with the required parameters.
- Stop the instance. 2.
- Create a minimum of two media interfaces and attach to the instance.

- Start the instance.
- Login to the virtual machine.
- 6. Reset the default password using either of the two ways:
 - Type the default password.
 - Use a SSH key.

If you type the default password, for a user account, it is acme followed by VM Id (without any space between acme and VM Id).

For e.g. acmed1b1f958-d8f0-4559-b3f1-60cf311b0d46



Obtain the VM Id from the JSON View of the VM in Azure.

For an admin account, the default password is packet followed by VM Id (without any space between packet and VM Id).

For e.g. packetd1b1f958-d8f0-4559-b3f1-60cf311b0d46

If you use a SSH key,

Change ssh-key file permissions to 400

```
ORACLE# chmod 400 <sshkey>
```

Type the below command:

```
ORACLE# ssh -i <sshkey> admin@<VM public IP>
```

- 7. Use the **setup product** command to setup the specific SBC product.
- 8. Use the **setup entitlement** command to self-configure the required entitlements.
- Use the interfaces mappingcommand to view the interfaces with their respective MAC addresses.
- Swap the interface showing incorrect MAC address with the interface having correct MAC address.



Verify the MAC address from Azure portal > VM > Network interface > Properties > MAC Address.

- **11.** Stop the instance.
- 12. Enable the accelerated networking for all media interfaces.
- 13. Start the instance.

Create and Deploy on Google Cloud Platform

You can deploy the SBC on Google Cloud Platform (GCP) in either standalone mode or high availability (HA) mode.

Oracle supports the following machine types.

Table 6-1 GCP Machine Types

Machine Type	vCPUs	Memory (GB)	vNICs	Egress Bandwidth (Gbps)	Max Tx/Rx queues per VM
n2-standard-4	4	16	4	10	4
n2-standard-8	8	32	8	16	8
n2-standard-16	16	64	8	32	16

Use the n2-standard-4 machine type if you're deploying an SBC that requires one management interface and only two or three media interfaces. Otherwise, use the n2-standard-8 or n2-standard-16 machine types for an SBC that requires one management interface and four media interfaces. Also use the n2-standard-8 or n2-standard-16 machine types if deploying the SBC in HA mode.

Before deploying your SBC, check the Available regions and zones to confirm that your region and zone support N2 shapes.

On GCP the SBC must use the **virtio** network interface card. The SBC will not work with the GVNIC

Prerequisites to Deploying an GCP Instance

Complete these steps before you deploy the SBC on Google Cloud platform.

Create a Virtual Private Cloud

On GCP, a virtual private cloud (VPC) must be created for each network interface. If you're deploying the SBC in standalone mode, create one management network and four media networks. If you're deploying the SBC in HA mode, create one management network, two HA networks, and four media networks.

- wancom0
- wancom1 (HA only)
- wancom2 (HA only)
- s0p0
- s1p0
- s0p1
- s1p1
- 1. From the Google Cloud navigation menu, click **VPC network** and then **VPC networks**.
- On the top bar, click CREATE VPC NETWORK.
- 3. Enter a name for this network.
- 4. Select the region.
- 5. Under **New subnet** define the subnet's name, region, and CIDR.
- 6. Select the IPv4 firewall rules for this network.

 If setting the firewall rules for the management interface wancom0, allow ICMP and SSH traffic. Otherwise, you can more clearly define firewall rules after creating the VPC.



Click CREATE.

Repeat this process for each network interface.

Create Firewall Rules

Each network interface needs a dedicated firewall rule to govern traffic on that interface. When you create a firewall rule, you specify what kinds of traffic the firewall should match (such as incoming TCP traffic on port 22 from 10.2.2.0/24) and whether the traffic is allowed or denied.

- From the navigation menu, click VPC network and then Firewall.
- 2. From the top menu, click CREATE FIREWALL RULE.
- Enter a name for this firewall rule.
- 4. Select the VPC network that this firewall rule will apply to.
- Set the Direction of traffic to ingress or egress.
- 6. Set the Action on match to allow or deny.
- Set the Targets field to All instances in the network.
- 8. Enter the source ranges that the firewall rule will match.
- 9. Select the protocols and ports that the firewall rule will match.
- 10. Click CREATE.

Repeat this process for each network interface.

The following is an example of the cloud shell command to create a firewall rule for the s0p0 media interface that allows outgoing TCP traffic on port 5060 to any destination.

```
gcloud compute --project=<PROJECT-NAME> firewall-rules create s0p0-fw-out \
    --direction=EGRESS \
    --priority=1000 \
    --network=sbc-s0p0 \
    --action=ALLOW \
    --rules=tcp:5060 \
    --destination-ranges=0.0.0.0/0
```

Enable Network Peering

Network peering allows internal IP addresses to connect to each other. The SBC uses network peering to route traffic between media interfaces.

- 1. From the navigation menu, click VPC network and then VPC network peering.
- 2. From the top menu, click CREATE PEERING CONNECTION and then Continue.
- 3. Enter a name for this peering connection.
- 4. Enter the name of the first VPC network.
- 5. Enter the name of the second VPC network.
- 6. Select both Import subnet routes with public IP and Export subnet routes with public
- 7. Click CREATE.

To create peering in both directions, repeat the process but swap the values of the "VPC network name" field and the "Your VPC network" field.

Create a Bucket

Create a bucket to contain the KVM disk image.

- From the navigation menu, click Cloud Storage and then Buckets.
- 2. Click CREATE.
- 3. Enter a name for this bucket.
- Click CREATE.

Upload the KVM Disk Image

Convert the KVM image to a raw disk and upload it to your Google Cloud Storage.

- 1. Sign in to the Oracle Software Delivery Cloud.
- Search for "Oracle Communications Session Border Controller", select release S-Cz10.0.0, and click Continue.
- 3. Under "Platforms/Languages," select All and click Continue.
- 4. Accept the License Agreement and click Continue.
- Download the KVM image.
- **6.** Unzip and untar the qcow2 file. For example:

```
unzip VXXXXXXX-01.zip
tar xf nnsCZ920-img-vm kvm.tgz
```

7. Convert the qcow2 file to a raw disk format.

```
qemu-img convert -f qcow2 -O raw nnSCZ920-img-vm kvm.qcow2 disk.raw
```



The filename must be disk.raw.

8. Compress the raw disk.

```
tar --format=oldgnu -Sczf compressed-nnSCZ920-image.tar.gz disk.raw
```

- 9. Upload the compressed raw disk to Google Cloud Storage.
 - **a.** From the Google Cloud navigation menu, select **Cloud Storage**, then **Buckets**, then the name of your bucket, and then **UPLOAD FILES**.
 - **b.** Select your tar file.
 - c. Click Open.

Create a Custom Image

Your custom image must be created with the Cloud Shell because the GUI does not provide a way to enable multiple subnets.

1. From the Google Cloud navigation menu, select **Compute Engine** and then **Images**.

- On the top bar, select CREATE IMAGE.
- 3. Enter a name for this image.
- Set the Source to Cloud Storage file.
- **5.** In the Cloud Storage file box, browse to the tar file you uploaded.
- 6. Click EQUIVALENT COMMAND LINE and then RUN IN CLOUD SHELL.
- 7. Append the following line to the command to enable multiple subnets.

```
--guest-os-features="MULTI_IP_SUBNET"
```

An example command, with whitespace introduced for clarity:

```
glcloud compute images <image-name> \
    --guest-os-features="MULTI_IP_SUBNET" \
    --source-uri=https://storage.googleapis.com/download/storage/v1/b/
<BUCKET_NAME>/o/<OBJECT_NAME>?alt=media
```

You now have a custom image you can use to create a virtual machine instance template.

Create an Instance Template

An instance template is a saved virtual machine configuration from which you can deploy a VM or a group of VMs.

- 1. From the navigation menu, click **Compute Engine** and then **Instance templates**.
- From the top menu, click CREATE INSTANCE TEMPLATE.
- Enter a name for the instance template.
- Set the Series field to N2 and set the Machine type field to one of the following:
 - n2-standard-4
 - n2-standard-8
 - n2-standard-16
- 5. Click CPU PLATFORM AND GPU to expand the CPU configuration.
- 6. Set CPU platform to Intel Ice Lake or later.
- Set the Boot disk to boot your custom image with a 40GB disk.
 - Under Boot disk, click CHANGE.
 - b. Select the **CUSTOM IMAGES** tab.
 - c. Set the Image field to the custom image you created.
 - d. Set the disk size to 40GB.
 - e. Click SELECT to close the Boot disk.
- Click Advanced options and then Networking.
- Set the Network interface card to VirtlO.
- 10. Add network interfaces.
 - If using a load balancer, add s0p0 as the first network interface.
 - b. If not using a load balancer, add wancom0 as the first network interface.

- c. Click ADD NETWORK INTERFACE to add the additional network interfaces.
- 11. Add your SSH public key.
 - a. Click Security and then MANAGE ACCESS.
 - b. Click ADD ITEM.
 - c. Paste in your SSH public key.



The SBC does not currently support ed25519 SSH keys.

12. Click CREATE.

Deploy a Standalone SBC

After creating your custom template, follow these steps to deploy a standalone instance of the SBC on Google Cloud Platform.

- 1. From the navigation menu, click **Compute Engine** and then **Instance templates**.
- 2. Click on the instance template you previously created.
- 3. Click CREATE VM.
- Enter a name for the SBC.
- 5. Select the region and zone where you want to deploy this SBC.



If you select a region different from the region in which the VPCs were created, the web interface prevents you from deploying the VM.

6. Click CREATE.

After the VM has been created, use the interface-mapping command to map virtual interfaces to the SBC's management and media interfaces.

Deploy High Availability SBCs

On the Google Cloud platform, you must create a load balancer to deploy the SBC in HA mode.

Limitations

On Google Cloud platform, a load balancer can only serve media traffic to the nic0 interface. When deploying a VM in HA mode, make sure to set the first interface as s0p0. One consequence of this limitation is that secondary media interfaces (s0p1, s1p0, s1p1) can only be used for internal trafic.

Reserve Public IPs

Reserve two static IP addresses to use for the load balancer: an access-side IP address and a core-side IP address.

- 1. From the navigation menu, click **VPC network** and then **IP addresses**.
- From the top menu, click RESERVE EXTERNAL STATIC ADDRESS.
- 3. Enter a name for this reserved static IP address. For example, access-ext-ip.
- 4. Select whether to reserve an IPv4 or IPv6 address.
- Set the **Region** to the region you decided to use.
- Leave Attached to set to None.
- Click RESERVE.

Repeat these steps to reserve a public IP address for the core side.

Reserve Static Internal IPs

In HA environments, each media interface needs a static IP address.

- 1. From the navigation menu, click VPC network and then VPC networks.
- 2. Select one of the media interfaces you previously created.
- Click the STATIC INTERNAL IP ADDRESSES tab and then RESERVE STATIC ADDRESS.
- Enter a name for the internal static IP. For example, s0p0-staticip.
- 5. Under Static IP address, select **Let me choose** and enter the IP address.
- 6. Set Purpose to Shared.
- Click RESERVE.

Repeat this process for each media interface.

Create a Health Check

To minimize the disruption due to a failover, create a health check that automatically switches from the active SBC to the standby SBC.

- 1. From the navigation menu, click Compute Engine and then Health checks.
- 2. From the top bar, select **CREATE HEALTH CHECK**.
- 3. Enter a name for this health check.
- 4. Under Scope, select Regional.
- 5. Select a port, such as 8888, on which to conduct health checks.
- Under Health criteria, set both Check interval and Timeout to 1 second.
- Click CREATE.

Create Instance Group

After you create an instance group, you can bring up two SBC instances at the same time.

- 1. From the navigation menu, click Compute Engine and then Instance templates.
- 2. Select your previously created instance template and click CREATE INSTANCE GROUP.
- 3. In the left column, confirm New managed instance group (stateless) is highlighted.
- Create a name for this instance group.



- Under Location, select the region and zone.
- 6. Under Autoscaling, set Autoscaling mode to Off.
- 7. Set the Minimum number of instances field to 2.
- 8. Click CREATE.

Create an External Load Balancer

Set up an external load balancer for the nic0 / s0p0 interface.

If the **Network services** panel does not appear in the navigation menu, expand the **MORE PRODUCTS** menu and scroll down to the networking section.

- 1. From the navigation menu, click **Network services** and then **Load balancing**.
- 2. Click CREATE LOAD BALANCER.
- 3. Under UDP Load Balancing, click START CONFIGURATION.
- 4. Select the following options:
 - From Internet to my VMs
 - Backend Service
- 5. Click CONTINUE.
- Enter a name for this load balancer.
- Under Backend configuration:
 - a. Set the **Instance group** to your previously created instance group.
 - **b.** Set the **Health check** to your previously created health check.
 - c. Click DONE.
- 8. Under Frontend configuration:
 - a. Define a name for this configuration.
 - b. Under IP address, select the previously reserved public IP address.
 - c. Under Ports, select All.
 - d. Click DONE.
- Click CREATE.

Create an Internal Load Balancer

Set up an internal load balancer to route private IP traffic on nic1, nic2, etc. interfaces.

If the **Network services** panel does not appear in the navigation menu, expand the **MORE PRODUCTS** menu and scroll down to the networking section.

- 1. From the navigation menu, click **Network services** and then **Load balancing**.
- 2. Click CREATE LOAD BALANCER.
- Under UDP Load Balancing, click START CONFIGURATION.
- Select Only between my VMs.
- 5. Click CONTINUE.
- 6. Enter a name for this load balancer.
- 7. Set the region to the same region you previously decided to use.



- 8. Set the **Network** to your s0p0 network.
- 9. Under Backend configuration:
 - a. Set the **Instance group** to your previously created instance group.
 - b. Set the **Health check** to your previously created health check.
 - c. Click DONE.
- 10. Under Frontend configuration:
 - a. Define a name for this configuration.
 - b. Under **Subnetwork**, select the previously created subnet for this network.
 - Under Internal IP, select Shared and set the IP address to your previously created reserved static IP.
 - d. Under Ports, select All.
 - e. Click DONE.

11. Click CREATE.

Repeat this process for each media interface.

Create Firewall Rules for Health Checks

Create a firewall rule to allow traffic from the load balancer to the SBCs. When you create a firewall rule, you specify what kinds of traffic the firewall should match (such as incoming TCP traffic on port 22 from 10.2.2.0/24) and whether the traffic is allowed or denied.

- 1. From the navigation menu, click VPC network and then Firewall.
- From the top menu, click CREATE FIREWALL RULE.
- 3. Enter a name for this firewall rule.
- Set the VPC network to the s0p0 network.
- Set the Direction of traffic to ingress.
- Set the Action on match to allow.
- Set the Targets field to All instances in the network.
- 8. Enter the source ranges that the firewall rule will match.
 This is the external IP address of the load balancer.
- Select the protocols and ports that the firewall rule will match.For example, if you previously set the health port to 8888, set the firewall rule to match traffic on port 8888.
- 10. Click CREATE.

After the two HA SBCs has been created, use the interface-mapping command to map virtual interfaces to the SBC's management and media interfaces. Make sure to map nic0 to s0p0.

Configure an HA Pair on GCP

Below is a sample configuration of an HA pair on GCP.

In this example, the IPs 104.196.215.97 and 35.211.156.206 are configured as the front end IP addresses of the external load balancer. The IPs 10.79.3.79 and 10.79.4.79 are configured as the front end IP addresses of the internal load balancer. These IPs can be added with the addhip-ip command in the network-interface configuration element.



Under **system-config**, enable the **use-sibling-core-datapath** parameter to utilize all virtual CPUs.

http-server	
name	s0p0 GCPHealthCheck FEIP1
realm	access
ip-address	104.196.215.97
http-port	8888
http-server	
name	s0p0_GCPHealthCheck_FEIP2
realm	access
ip-address	35.211.156.206
http-port	8888
http-server	
name	s1p0 GCPHealthCheck
realm	core
ip-address	10.79.4.79
http-port	8888
http-server	
name	s0p0 GCPHealthCheck
realm	access
ip-address	10.79.3.79
http-port	8888
local-policy	
from-address	3.90.177.223
to-address	104.196.215.97
policy-attribute	101.130.210.37
next-hop	10.79.7.100
realm	core
local-policy	COIC
from-address	3.90.177.223
to-address	35.211.156.206
policy-attribute	33.211.130.200
next-hop	10.79.7.100
realm	core
media-manager	COIC
network-interface	
name	s0p0
ip-address	10.79.3.79
pri-utility-addr	10.79.3.94
sec-utility-addr	10.79.3.95
netmask	255.255.255.0
gateway	10.79.3.1
hip-ip-list	10.79.3.79
111p 1p 110c	104.196.215.97
	35.211.156.206
icmp-address	10.79.3.79
network-interface	10.79.3.75
name	s1p0
ip-address	10.79.4.79
pri-utility-addr	10.79.4.79
sec-utility-addr	10.79.4.92
netmask	255.255.255.0
	10.79.4.1
gateway hip-ip-list	10.79.4.79
icmp-address	
TOMP-address	10.79.4.79

network-interfa	ice			
name			wancom1	
pri-uti	lity-add	r	10.79.1	.61
sec-uti	lity-add	r	10.79.1	.62
netmask			255.255	.255.0
phy-interface				
name			s0p0	
operati	on-type		Media	
phy-interface				
name			s1p0	
operati	on-type		Media	
slot			1	
phy-interface				
name			wancom1	
port			1	
duplex-	mode			
speed				
-	health-s	core	8	
realm-config				
identif	ier		access	
network	-interfa	ces	s0p0:0.4	4
realm-config			21,7111	
identif	ier		access1	
addr-pr			35.211.3	156.206
=	-interfa	ces	s0p0:0	200,200
realm-config			2010.0	
identif	ier		access2	
addr-pr			104.196	215 97
	-interfa	ces	s0p0:0	
realm-config	IIICCIIG	000	5050.0	
identif	ier		core	
	-interfa	CAS	s1p0:0.4	4
redundancy-conf			Sipo.o.	-
peer peer	. 19			
PCCI	name			ha-sbc0
	type			Primary
	destina	tion		rrimary
		address		
10.79.1.61:9090		addless		
10.79.1.01.9090		network-interface		wancom1:0
2002		network-interrace		wallColli. C
peer	2222			ha aha1
	name			ha-sbc1
	type destina	+		Secondary
	destina			
10 70 1 60 000		address		
10.79.1.62:9090				1 (
		network-interface		wancom1:0
session-router				
sip-config				
sip-interface	,		1	
realm-i			access1	
sip-por				25 211 156 226
	address			35.211.156.206
sip-interface	1		_	
realm-i			access2	
sip-por	ìt.			

104.196.215.97 address sip-interface ${\tt realm-id}$ core sip-port address 10.79.4.79 steering-pool 10.79.3.79 ip-address start-port 20000 end-port 60000 realm-id access steering-pool ip-address 10.79.4.79 start-port 20000 end-port 60000 realm-id core steering-pool 104.196.215.97 ip-address 20000 start-port end-port 60000 realm-id access2 steering-pool 35.211.156.206 ip-address start-port 20000 end-port 60000 realm-id access1 system-config transcoding-cores enabled use-sibling-core-datapath

7

Boot Management

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

Boot management consists of tasks working with the following:

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

Boot Loader Overview

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

Stage3 Boot Loader

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

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

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

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

Stage 3 boot loader upgrade procedure can be found in the Update the Stage 3 Bootloader section of this guide.



The SBC does not support uploading the boot loader by way of the Web GUI.

Boot Parameters

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

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

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

Note:

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

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

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

The bootparam configuration list is shown below.



```
Console Baudrate : 115200
Other
[Acme Boot]: ?
                       - print this list
9
                       - boot (load and go)
                       - print boot params
р
С
                       - change boot params
                       - print boot logo with version
V
                       - reboot
r
                       - show license information
 S
```

Boot Parameter Definitions

The system displays all boot parameters when you configure them after a boot interrupt. The system hides some boot parameters from the ACLI so that you do not attempt to configure them. If changed improperly, these parameters can cause the system to sop responding.

The following table defines each of the parameters that the system displays when you perform configuration 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.



VLANs over management interfaces are supported only on the Acme Packet 1100.

Note:

The acquire-config command is not supported on management interfaces that use both VLANs and IPv6.

Netmask	Netmask portion of the wancom0 IP Address.
Gateway	Network gateway that this wancom0 interface uses.
IPv6 address	Version 6 IP address/mask of wancom0. Configure the mask as a forslash (/) after the address followed by the mask in number of bits.
IPv6 Gateway	Version 6 network gateway that this wancom0 interface uses.
Host IP	IP Address of FTP server from which to download and execute a software image.
FTP Username	FTP server username
FTP password	FTP server password



Boot Parameter	Description
Flags	Codes that signal the system from where to boot. Also signals the system about which file to use in the booting process. This sequence always starts with 0x (these flags are hexadecimal).
Target Name	Name of the Oracle Communications Session 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 must be unique among Oracle Communications 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:

- 0x0000001 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.





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

Change Boot Parameters from the ACLI

To access and change boot parameters from the ACLI:

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

```
ORACLE# configure terminal
```

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

```
ORACLE(configure) # bootparam
'.' = clear field; '-' = go to previous field; q = quit
Boot File : /boot/nnScz100.bz
```

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

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

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

```
ORACLE(configure) # bootparam
'.' = clear field; '-' = go to previous field; q = quit
Boot File : /boot/nnSCz100.bz /boot/nnSCz200.bz
```

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

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

```
ORACLE (configure) #
```

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

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



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

Change Boot Parameters by Interrupting a Boot in Progress

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

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

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

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

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

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

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

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

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

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

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

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

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



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

Displaying Files from Boot and Images directory

You can view the list of image files (.bz files) from **/boot** and **/code/images** directories using the I (small letter L) option from the boot menu prompt. This allows you to select an alternate image file when the boot file gets corrupted. This option is available only when FIPS mode is disabled. You can view the files in a page view when there are more than 15 files in the directories. When SBC boots over network using ftp/sftp with the selected image file, this option does not update the boot param. This option displays image files from both **/boot** and **/code/images** directories irrespective of R226 license.

To use this option from the boot menu:

```
[Acme Boot]: p
Boot File : /boot/bzImage
IP Address
                : 172.44.12.89
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 : 0x0000
Target Name : ORACLE
                : 0x00000040
Flags
Console Device : COM1
Console Baudrate: 115200
Other
[Acme Boot]: ?
 ?
                       - print this list
 @
                       - boot (load and go)
                       - print boot params
 р
                        - change boot params
 С
                        - Oracle Rescue Access sub-menu
 0
                        - print boot logo with version
 ٦7
 r
                        - list diagnostic images
 d
                        - show license information
 S
 1
                        - show boot images
Boot flags:
   0x02 - enable kernel debug
   0x04 - disable crashdumps and enable minidump
   0x10 - enable debug login
   0x40 - use DHCP for wancom0
         - use TFTP instead of FTP
   0x80
```

```
[Acme Boot]: 1
1:
     /boot/nnSCZ900p2.bz
2:
     /boot/nnSCZ900p1.bz
     /boot/nnSCZ840p2.bz
3:
4:
     /boot/a.bz
     /boot/b.bz
5:
6:
     /boot/c.bz
     /boot/d.bz
7:
8:
     /boot/e.bz
     /boot/f.bz
9:
     /code/images/nnSCZ900p2.bz
10:
11:
       /code/images/nnSCZ900p1.bz
12:
     /code/images/nnSCZ840p2.bz
13:
      /code/images/a.bz
14:
       /code/images/b.bz
15:
       /code/images/c.bz
Few more entries press c to continue or select from list to boot or any other
key to quit listing...:c
[Acme Boot]:[Boot Image]: 3
```

This option is unsupported on OVM platform.

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.



For systems without a separate boot flash, the **format system-disk** command combines the **lopt** and **lopt/crash** into a single system partition. This includes all systems other than the Acme Packet 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 8-1 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 8-2 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 8-3 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	/mnt/ <user-label></user-label>	user-defined percentage of remaining space

Note:

Oracle recommends creating a single mount point for data partitions, such as /mnt/app, and then using subfolders for specific purposes, such as /mnt/app/HDR or /mnt/app/CDR.

A

Caution:

Creating a folder directly under /mnt without first formatting a partition is not supported and likely to result in data loss. Use the format command to create mount points.

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:
                   8000000 bytes
1: /opt
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.



```
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.



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.



9

Interface Considerations for VM Platforms

Oracle Communications 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 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 Session Border Controller **phy-interface** elements, which is documented in the *ACLI Configuration Guide*.

Media Interface Consideration

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, Use either Paravirtual or SR-IOV I/O mode for media interfaces.

COTS Interfaces

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

COTS Network Interfaces

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

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

The interface-mapping Branch

The **interface-mapping** branch resides at the root level of the Oracle Communications Session 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 Session Border Controller application's configuration. The system accomplishes this using physical interface MAC addresses and ACLI configuration interface naming.

Sample default Oracle Communications 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
      00:50:88:BC:21:FF
s1p0
                             #generic
s0p1
      00:50:88:BC:41:A2
                             #generic
s1p1
      00:50:88:BC:31:AC
                             #generic
s0p2
     FF:FF:FF:FF:FF
                             #dummy
s1p2
      FF:FF:FF:FF:FF
                             #dummy
s0p3
       FF:FF:FF:FF:FF
                             #dummy
s1p3
       FF:FF:FF:FF:FF
                             #dummy
```

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

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

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

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

Working with the interface-mapping branch

Interface identification on the Oracle Communications Session 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 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
s1p0	00:16:3E:30:00:2D	# Second media interrface

interface-mapping

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

Command	Description
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></seconds></ethernet>	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</ethernet>	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.</eth-if-name>
interface-mapping delete <ethernet if="" name=""></ethernet>	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.



Command	Description
interface-mapping swap <ethernet if="" name1=""> <ethernet if="" name2=""></ethernet></ethernet>	Swap the mapping of Ethernet interface names against the available MAC physical interfaces. For example, you can first execute the interfacemapping show command to display the current information. interface-mapping show
	wancom0 00:16:3E:30:00:2A # control port, onboard MAC
	wancom1 00:16:3E:30:00:2B # 2nd control port, onboard MAC
	s0p0 00:16:3E:30:00:2C # PCI left side
	s1p0 00:16:3E:30:00:2D # PCI right side
	Then you can execute the interface-mapping swap command.
	interface-mapping swap s0p0 s1p0
	wancom0 00:16:3E:30:00:2A # control port, onboard MAC
	wancom1 00:16:3E:30:00:2B # 2nd control port, onboard MAC
	s0p0 00:16:3E:30:00:2D # PCI right side
	s1p0 00:16:3E:30:00:2C # PCI left side
	A warning message appears. Once you confirm the action, the MAC addresses and their corresponding labels are swapped in the /boot/mactab/file.



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

Serial Interfaces

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

Access to the Oracle Communications Session 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.

Flash Drive Installation via Boot Media Creator

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

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

For example, nnSCZ930-img-usb.exe is the BMC distribution of software version S-CZ9.3.0.

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


Creating a Build Image

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

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


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

Figure 10-1 BMC - Step 3




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

Figure 10-2 BMC - Step 4




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

Figure 10-3 BMC - Step 5



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


Figure 10-4 BMC - Step 6



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


- Installation Type: Commission
- Console Port: Default
- Include Preload tarfile: unchecked
- Generate installation logs: unchecked
- Click Next. BMC confirms your settings.

Figure 10-5 BMC - Step 7




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


Figure 10-6 BMC - Step 8



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


Figure 10-7 BMC - Step 9





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

Figure 10-8 BMC - Step 10



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

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

Software Upgrade

This section provides information about how to upgrade your Oracle Communications 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 that what was delivered, making software upgrade a required step in platform preparation.

The Check Upgrade Readiness Command

The SBC provides you with the ability to examine system components and status to determine whether an upgrade can occur seamlessly. If not, the system provides you with information that describes changes you should make prior to attempting an upgrade. Review the output of this command three times during an upgrade cycle:

- 1. Run the **check-upgrade-readiness** command before any software upgrade. Examine the command output to identify and resolve any potential upgrade issues.
- 2. Also note that the system runs the check-upgrade-readiness command automatically upon the first reboot after a software upgrade. If the system finds significant issues after running the command, it may revert back to booting with the previous software version.
- Run the check-upgrade-readiness after upgrade to check for any potential operational issues.

Note that the **verify-config** command can also provide important system status information that can assist you during an upgrade.

Upgrade Checklist

Before upgrading the Oracle Communications Session Border Controller software:

- 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 Oracle Communications Session Border Controller image file in the boot parameters.
- **3.** Run the **check-upgrade-readiness** command and examine its output for any recommendations or requirements prior to upgrade.
- 4. Verify the integrity of your configuration using the ACLI verify-config command.
- 5. Back up a well-working configuration. Name the file descriptively so you can fall back to this configuration easily.
- 6. Refer to the Oracle Communications Session Border Controller Release Notes for any caveats involving software upgrades.

Check Upgrade Readiness

The Oracle Communications Session Border Controller (SBC) includes the **check-upgrade-readiness** ACLI command, which presents system information arranged to clearly tell you if you need to perform any tasks before you upgrade.

The **check-upgrade-readiness** command performs system diagnostics and configuration checks to generate a system status report. The command provides you with compatibility information with respect to new software version and boot loader version, boot parameter validity checks, NIU version, FPGA versions and so forth. In addition, the command displays system status information, such as CPU and memory utilization, that could impact the upgrade.

For HA deployments, the **check-upgrade-readiness** command provides a warning to the user if they are performing the command on the active and displays the application processes that are not synchronized.

Oracle recommends examining the output of this command at least twice:

- 1. Prior to upgrade—The command can display potential risks with the upgrade.
- 2. After upgrade—The system runs the command automatically when booting to a new software version to identify any component version incompatibilities. If the system finds problems, it stops the boot process and reboots to the prior version.

The command has a brief and a verbose output. The abbreviated output shows basic information and lists any issues.

```
ORACLE#check-upgrade-readiness
Warning: Some of the functional modules are overloaded, upgrading system in
such overload state may lead to instabilities. Details along with current
running version as follows...
Version info:
       Product name : Acme Packet 4600
       Release version: TSCZ0.0.0 Cycle 40 (WS Build 519)
       Platform : NN4600
Validate platform parameters:
       WARNING : session agents (9/10)
Validate CPU usage:
               : CPU Core 01 (Percent Usage:96)
       ERROR
Validate Standby HA role:
       WARNING : System with Active HA role. Upgrade standby system
first.
Check alarm status:
       WARNING : 1 unclear alarms exists in the system.
```

The verbose output shows the results of each check.

```
ORACLE#check-upgrade-readiness verbose
Warning: Some of the functional modules are overloaded, upgrading system in such overload state may lead to instabilities. Details along with current running version as follows...

System version info:

Product name : Acme Packet 4600

Release version : TSCZ0.0.0 Cycle 40 (WS Build 519)

Platform : NN4600
```



```
Validate Memory Usage:
      OK : Memory usage
                                           (Percent Free: 93)
Validate CPU usage:
       OK : CPU Usage
OK : CPU Core 00
                                          (Percent Usage:77)
                                           (Percent Usage:89)
       WARNING : CPU Core 01
                                           (Percent Usage:93)
       OK : CPU Core 02
                                           (Percent Usage:44)
       OK
             : CPU Core 03
                                            (Percent Usage:81)
Validate Active and Standby sync status:
       OK : System not in HA pair.
Validate Standby HA role:
       WARNING: System with Active HA role. Upgrade standby system first.
Validate bootparams:
       OK : IP address
                                             (10.196.131.112)
       OK
             : netmask
                                             (255.255.224.0)
       OK
             : gateway
                                             (10.196.128.1)
             : host ip
       OK
                                             (10.196.0.83)
             : IPV6 address
       OK
                                             ()
             : IPV6 gateway
                                             ()
       OK
             : space in /boot
                                             (Percent Free: 84)
Validate platform parameters:
                    : sessions
                                             (0/32000)
                     : Sip Rec Sessions
       OK
                                             (0/16000)
                     : SRTP Sessions
       OK
                                             (0/1000)
       OK
                     : MSRP Sessions
                                             (0/2000)
       OK
                     : Licensed Capacit
                                             (0/32000)
       OK
                      : HMU flows
                                             (0/64000)
       WARNING : session agents
                                             (9/10)
       OK
                     : ARPs
                                             (0/4120)
                    : INTFC Flows
: Trusted Entries
       OK
                                             (0/4096)
       OK
                                             (3/8192)
                     : Untrusted Entries
                                             (1/4096)
       OK
                     : Media Entries
                                             (1/64000)
       OK
                     : Deny Entries
                                             (0/32768)
                     : Internal Flows
       OK
                                             (2/34088)
       OK
                     : RFC 2833 Flows
                                             (0/32000)
       OK
                     : QoS Sessions
                                             (0/64000)
                     : Xcoded Sessions
       OK
                                             (0/20000)
       OK
                     : Datapath TCP Sockets (0/128000)
                     : Datapath TLS Sockets (0/512)
       OK
Check alarm status:
       WARNING: 2 unclear alarms exists in the system.
2 alarms to show
ID Task Severity First Occurred Last Occurred 327691 94 4 2018-04-05 11:27:29 2018-04-05 11:27:29
                                                  Last Occurred
Count Description
2 all cdr push receivers have failed
```

```
131161 3350 4 2018-04-05 17:47:24 2018-04-05 17:47:24

Count Description

1 CPU core0 is at 97 percent, over major threshold of 90 percent done
```

Download the Software

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

Before You Begin

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

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

- 1. Go to Oracle Cloud Software Delivery at https://edelivery.oracle.com/, and sign in.
- On the Terms and Restrictions page, in the Oracle Trial License Agreement section, select one of the following:
 - If you want a trial license, select Yes.
 - If you have a license, select Or.
- 3. On the Terms and Restrictions page, in the Export Restrictions section, select Yes.
- Click Continue.
- 5. On the Media Pack Search page, do the following:
 - Select a Product Pack
 - Select a Platform
- 6. Click Go.

The system displays a list of software for the selected product pack and platform.

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

Next Steps

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

Check /boot for free space

Perform this procedure for stand-alone and HA deployments, ensuring adequate space on all applicable machines.

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

```
SBC# 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**.

Select Upgrade Procedure

Customers have two options for upgrading the SBC.

- If you are upgrading from one major or minor release to a different major or minor release, upload a new bootloader and boot file.
 For example, if upgrading from release S-Cz8.4.0 to S-Cz9.0.0 (a new major release), upload a new bootloader and boot file. Or if upgrading from release S-Cz9.2.0 to S-Cz9.3.0 (a new minor release), upload a new bootloader and boot file. See "Update the Stage 3 Bootloader" for details.
- 2. If you are upgrading to the latest patch release within your Major. Minor release stream, upload your boot file and update the boot parameters.
 For example, if upgrading from S-Cz9.1.0p5 to S-Cz9.1.0p7, upload your boot file and update the boot parameters. See "Stand-alone Upgrade" for details.

Upload the Stage 3 Boot Loader and System Image

While the boot loader is generally backward compatible with previous releases, Oracle recommends that you install a boot loader from the same Major. Minor release as the system image. Installing a maintenance or patch release within a Major. Minor release does not normally require updating the boot loader. Perform this procedure for stand-alone and HA deployments, ensuring you are bootingall applicable machines using the correct bootlaoder.

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 /code/images/ before you update the boot parameters to use the new software image file. The boot loader file must be renamed using the **set-boot-loader** command on the target system. When upgrading an HA pair, you must perform the upgrade procedure on each HA node. This includes rebooting the systems after setting the bootloader and image. When performing this task, make sure to upgrade and reboot the standby system first.

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

1. SFTP the software image (*.bz) and Stage3 boot loader (*.boot) to /code/images/.



```
Uploading nnSCZ900p1.boot to /code/images/nnSCZ900p1.boot nnSCZ840p3.boot 100% 15MB 30.8MB/s 00:00 Uploading nnSCZ840p3.bz to /code/images/nnSCZ900p1.bz nnSCZ900p1.bz 100% 145MB 48.3MB/s 00:02 sftp> bye [Downloads]$
```

- 2. SSH to your target machine.
- 3. Run **set-boot-loader** with the path to the new bootloader.

```
ORACLE# set-boot-loader /code/images/nnSCZ900p1.boot
Verifying signature of /code/images/nnSCZ840p3.boot
Version: Acme Packet SCZ9.0.0 Patch 3 (Build 188) 202010201742

Image integrity verification passed

Successfully copied /code/images/nnSCZ900p1.boot to /boot/bootloader
ORACLE#
```

4. Run **set-boot-file** with the path to the new software image.

```
ORACLE# set-boot-file /code/images/nnSCZ900p1.bz
Verifying signature of /code/images/nnSCZ840p3.bz
Version: Acme Packet SCZ9.0.0 Patch 3 (Build 188) 202010201720

Image integrity verification passed
old boot file /boot/bzImage being replaced with /code/images/nnSCZ900p1.bz
ORACLE#
```

5. Reboot.

Stand-alone Upgrade



A stand-alone upgrade incurs system downtime. Plan your upgrade accordingly.

The following procedure describes how to upgrade an SBC with a new software image.

- SFTP your new boot image to the /boot directory of the SBC.
 If your /boot directory uses restricted permissions, upload to the /code/images/directory.
- 2. Change the boot configuration parameters to use the new image.



Changing the **Boot File** boot parameter configuration performs the same function as running the **set-boot-file** command.

In the ACLI configure terminal menu, type **bootparam** and press <Enter> to display the list of boot parameters. Stop when you reach the **Boot File** boot parameter and type the appropriate file name next to the previous file name.

In the following example, /boot/nnSCZ920p2.bz is the path of the old release and /boot/nnSCZ920p4.bz is the path of the new release.

```
SBC1# configure terminal
SBC1(configure)# bootparam

'.' = clear field; '-' = go to previous field; q = quit

Boot File : /boot/nnSCZ920p2.bz /boot/nnSCZ920p4.bz
```

- 3. Press <Enter> to continue scrolling through the boot parameters.
- **4.** Reboot the SBC using the **reboot** command.

The SBC 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.



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

Figure 11-1 Configure and Reboot Standby First for HA Procedures





Note:

Do not upgrade a non-LI HA deployment with an LI image or vice-versa if you wish to perform a hitless upgrade. This procedure results in a non-hitless upgrade, requiring that you reboot devices per your upgrade procedure, and then reboot all upgraded devices again to establish the new deployment type.

HA Upgrade Procedure

This procedure upgrades HA deployments.

Note:

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

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

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

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

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

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

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

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

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

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



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

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

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

Note:

Changing the **file name** boot parameter configuration, as shown below, performs the same function as running the **set-boot-file** command.

Note:

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

Access the boot parameters on ORACLE-2:

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

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

The following example uses the filenames /boot/nnSCZ840m5.64.bz and /boot/nnSCZ900.64.bz.

```
ORACLE-2# configure terminal

ORACLE-2(configure)# bootparam

'.' = clear field; '-' = go to previous field; ^D = quit

boot device : eth0

processor number : 0

host name : boothost

file name : /boot/nnSCZ840m5.64.bz /boot/nnSCZ900.64.bz
```

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

Reboot ORACLE-2.

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

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



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

```
ORACLE-2# show version
Acme Packet 4600 SCZ9.0.0
Build Date=09/09/15
```

- 7. Use the **show health** command to confirm that ORACLE-2 is the standby system.
- As you did for ORACLE-2, upload the SBC software image file and stage three bootloader to the /boot directory using an SFTP client. (See the instructions on updating the Stage 3 Bootloader.)
- Trigger a switchover from ORACLE-1 so that the standby system transitions to active, and vice-versa.

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

- **10.** Wait while ORACLE-2 transitions to the active state, then confirm that ORACLE-1 and ORACLE-2 are fully synchronized as explained in step 1.
- 11. Reboot the newly-standby ORACLE-1.

```
ORACLE-1# reboot
```

12. As you did for ORACLE-2, configure the boot parameters on ORACLE-1 to boot from the new software image. Then reboot ORACLE-1.

```
ORACLE-1# reboot

WARNING: you are about to reboot this SD!

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

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

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



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

HA Backout Procedure

If you reach the point in your upgrade procedure where you have upgraded both Oracle Communications Session 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.



Note:

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

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

 Change the boot parameters on ORACLE-1 to use the appropriate Release S-CZ9.0.0 software image.

Note:

Changing the **file name** boot parameter configuration, as shown below, performs the same function as running the **set-boot-file** command.

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

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

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

The following example uses the filenames /boot/nnSCZ900.64.bz and /boot/nnSCZ840.64.bz.

In the example above, type the appropriate Release S-CZ8.4.0 file name next to the Release S-CZ9.0.0 file name. Press <Enter> to continue scrolling through the boot parameters.

Exit to the main Superuser prompt.

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

2. Reboot ORACLE-1.



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

ORACLE-1# show version
Acme Packet 4600 SCZ9.0.0
Build Date=01/09/15

4. Initiate a switchover on ORACLE-1.

ORACLE-1# notify berpd force

At this point, ORACLE-1 becomes the active system running Release S-CZ8.4.0. ORACLE-2 is now the standby system running Release S-CZ9.0.0.

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



A

Physical Interfaces on Acme Packet Platforms

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

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

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

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

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

- SNMP
- SSH
- SFTP
- ACP/XML
- Logs sent from the Oracle Communications Session Border Controller
- The boot file used to boot the Oracle Communications 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	Required	Required




Parameter	Network Media Interface	Wancom1 and wancom2 Network Management Interfaces
operation-type	Required	Required
port	Required	Required
slot	Required	Required
virtual-mac	Optional	Invalid
admin-state	Required	Invalid
auto-negotiation	Required	Invalid
duplex-mode	Required	Invalid
speed	Required	Invalid
wancom-health-score	Invalid	Optional

Acme Packet 1100 Physical Interfaces

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

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

Figure A-1 Acme Packet 1100 - Rear View



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

- USB—For use only by Oracle personnel.
- SER MGT (Console)—Provides serial access for administrative and maintenance purposes.
- NET MGT (wancom0)—The system uses these 10/100/1000 Base-T Ethernet ports for device management functions and HA.
- INT—The system uses this 10/100 Mbps port for signaling and media traffic that is outbound with respect to the Remote Office/Branch Office (ROBO). This traffic originates from the ROBO.

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

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

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

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


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

Acme Packet 3900 Physical Interfaces

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

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

Figure A-2 Acme Packet 3900 - Rear View



Ports key

- T1/E1—For Time Division Multiplexing (TDM) quad span
- SER MGT—For console access for administrative and maintenance purposes
- MGMT0—For EMS control, RADIUS accounting, CLI management, SNMP queries and traps, and other network management functions
- MGMT1 and MGMT2—For High Availability (HA), or for network management with no HA configuration

- USB—For a storage device, or for installing software
- P0 P3—For signaling and media traffic on copper or fiber optic cable

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

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

Hardware Support

The Acme Packet 3900 hardware provides the following:

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

SNMP Hardware Reporting

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

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

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

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

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

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

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

Acme Packet 3900 MIBS Paths

Paths for Acme Packet 3900 MIBS.

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

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

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

```
SNMPv2-SMI::mib-2.47.1.1.1.1.2.4 = STRING: "500 Watt Power Supply"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.5 = STRING: "Assy, 2-fan unit of 40x10"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.6 = STRING: "Sensor of fan speed"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.7 = STRING: "Assy, Acme Packet 3900 Main Board"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.8 = STRING: "Sensor of temperature"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.9 = STRING: "Management Port 0 10/100 Ethernet Copper"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.10 = STRING: "Management Port 1 10/100 Ethernet Copper"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.11 = STRING: "Management Port 2 10/100 Ethernet Copper"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.12 = STRING: "Media port - Logical Slot 0 Port 0"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.13 = STRING: "Media port - Logical Slot 0 Port 1"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.14 = STRING: "Media port - Logical Slot 1 Port 0"
```

Acme Packet 3950/4900 Physical Interfaces

The Acme Packet 3950/4900 mainboard contains all the management and media ports, except for the optional T1/E1 ports. When installed, the T1/E1 ports are on a separate network interface card.

The following illustration shows the rear panel, which you need to know about when you perform the phy-interface configuration.

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

P1 P3 T1/E1 P2 **P4**

SER MGT

MGT0

MGT1

MGT2

Figure A-3 Acme Packet 4900 - Rear View

Ports key

- T1/E1—For Time Division Multiplexing (TDM) quad span
- SER MGT—For console access for administrative and maintenance purposes
- MGMT0—For EMS control, RADIUS accounting, CLI management, SNMP queries and traps, and other network management functions
- MGMT1 and MGMT2—For High Availability (HA), or for network management with no HA configuration

- USB—For a storage device, or for installing software
- P0 P3—For signaling and media traffic on copper or fiber optic cable at 1Gbps
- P4 P5—For signaling and media traffic on fiber optic cable at 10Gbps

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	0	2	
P3	Media	0	3	
P4	Media	0	4	
P5	Media	0	5	
T1/E1	Media	2	0	

Hardware Support

The Acme Packet 4900 hardware provides the following:


- 1 management interface at 1Gbps
- 2 HA interfaces at 1Gbs
- 4 media and signalling interfaces at 1Gbs
- 2 media and signalling interfaces at 10Gbs
- 2 USB ports
- (Optional) 1 quad-span Time Division Multiplexing (TDM) PCIe card

Acme Packet 4600 Physical Interfaces

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

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

Figure A-4 Acme Packet 4600 - Rear View



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

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



WARNING:

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

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

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


Acme Packet 6100 Physical Interfaces

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

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



Figure A-5 Acme Packet 6100 - Rear View



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

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

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

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

Acme Packet 6300/6350 Physical Interfaces

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

Acme Packet 6300/6350 Management Interfaces

Management interfaces include $3 \times 10/100/1000$ Ethernet interfaces labeled Mgmt 0, 1, and 2. (These interfaces are often referred to as wancom0, 1, and 2.) Mgmt 0, is for SSH access to the ACLI. The other two interfaces are used for state replication for High Availability (HA). For HA, connect these interfaces directly using a crossover cable.

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

The USB port is reserved for use by Oracle personnel, only. The following illustration displays the rear view of the Acme Packet 6300 and 6350 platform and labels the interfaces per the following table, which identifies the interface information per the software.

Figure A-6 Acme Packet 6300/6350 Management Interfaces



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

Graphic Label	NIU Label	Operation-type	Slot	Port	
1	Console	NA STATE	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	

Signaling and Media Interfaces

The signaling and media interfaces provide network connectivity for signaling and media traffic. Each interface can connect to a network at 10-Gigabit speeds.

Network Interface Units (NIUs) are available in the following configurations:

- 2-port 10GbE NIU
- 4-port 10GbE NIU

The optical 10GbE cards can accept an LC fiber connector using either single mode or multimode cable.

Mixed transceiver types are not supported on SFP+-based NIUs because compliance testing shows that the NIU SFP+ ports must be populated with identical SFP+ types.

The Acme Packet 6350 supports the 4x10GbE NIU to provide greater session scaling and Packet Processing Module (PPM) support. The 4x10GbE NIU contains two Cavium 78xx chips running in parallel, and each one is attached to two 10GbE media interfaces. Each chip contains 64GB of DDR4 RAM. The increase in DDR4 RAM enables greater scaling capacity, and the use of DDR4 running at 1050MHz provides a memory performance boost.



Note:

When using the 4x10GbE NIU, you must put it in slot 0.

Acme Packet 6300

Acme Packet 6300 Media Interface

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

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

Note:

Do not put any 2x10Gb NIU into slot 2.

Figure A-7 Acme Packet 6300 Media Interfaces



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

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



Acme Packet 6350

Acme Packet 6350 Media Interface

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

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



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

The Acme Packet 6350 supports up to 4x10GbE ports each running at full duplex line rate for all packet sizes. The following illustration shows the media ports numbering pattern.




Figure A-8 Acme Packet 6350 - Media Interface

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

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

Dual Port and Quad Port NIUs

The Acme Packet 6350 supports Dual 10 GbE and Quad 10 GbE Network Interface Units (NIU). The Dual NIU contains two 10G interfaces, while the Quad NIU contains four 10G interfaces to provide greater session scaling capacity and Packet Processing Module (PPM)

support. The Quad NIU also includes an internal Ethernet switch to allow for more flexible traffic loading to the multi-core processor.

The system supports using either the Dual NIU or the Quad NIU, but not at the same time. If you want to replace the Dual NIU with the Quad NIU, do the following

- Acme Packet 6350 owners using the Dual 10GbE NIU, can replace the Dual NIU with the Quad NIU.
- Acme Packet 6300 owners must purchase the Acme Packet 6350 and specify the Quad 10 GbE NIU.

If you own the Acme Packet 6350 and you replace the Dual 10 GbE NIU with the Quad 10 GbE NIU, you must re-set the port numbers in your PHY Interfaces, if you configured them by slot and port. The order of the port numbers for the Quad 10 GbE NIU differs from the Dual 10 GbE NIU, as follows.



Dual 10GbE NIU Port Numbering



Quad 10GbE NIU Port Numbering







The Quad 10 GbE NIU must go in slot 0.

Acme Packet 6400 Physical Interfaces

The Acme Packet 6400 integrates interfaces for console and network management into the rear of the chassis. The system also includes four media interfaces.

Acme Packet 6400 Management Interfaces

Management interfaces include $3 \times 10/100/1000$ Ethernet interfaces labeled MGT 0, MGT 1, and MGT 2. (These interfaces are often referred to as wancom0, wancom1, and wancom2.) MGT 0 is for SSH access to the ACLI; the other two management interfaces are used for high availability (HA). For HA, connect these interfaces directly using a crossover cable.

The SER MGT port is the serial console port.

Figure A-9 Acme Packet 6400 Management Interfaces



Signaling and Media Interfaces

The signaling and media interfaces provide network connectivity for signaling and media traffic. Each interface can connect to a network at 10-Gigabit speeds through the SFP+ connector.

The optical 10GbE cards can accept an LC fiber connector using either single mode or multimode cable. Mixed transceiver types are not supported on SFP+ ports because compliance testing shows that the SFP+ ports must be populated with identical SFP+ types. The Acme Packet 6400 supports the 4x10GbE network interface ports to provide greater session scaling and Packet Processing Module (PPM) support.

The following table maps the connection labels on the rear interfaces to the **operation-type**, **slot** and **port** attributes in the **phy-interface** configuration element.

Network Interface Label	operation-type	slot	port	
MGT 0	Control	0	0	
MGT 1	Control	0	1	
MGT 2	Control	0	2	
SER MGT	NA	NA	NA	
P0	Media	0	0	
P1	Media	0	1	
P2	Media	0	2	
P3	Media	0	3	



B

VNF Metadata and Userdata Example

The Heat templates provided by Oracle support the Newton and Pike releases. The following information may be useful for users who wish to implement their own orchestration template files for a different Openstack release.

For information about the OpenStack metadata and userdata files, see the OpenStack documentation.



When using the Heat template with the Newton or Pike releases, customers do not create, modify, delete, or interact with the OpenStack metadata and userdata files.

Metadata JSON Schema

```
"meta": {
    "bootparams.console": {
        "type": "string"
    "bootparams.gateway": {
        "type": "string"
    "bootparams.hostname": {
        "type": "string"
    "bootparams.ip0": {
        "type": "string"
    "bootparams.ip6": {
        "type": "string"
    },
    "bootparams.ip6gateway": {
        "type": "string"
    "bootparams.netmask0": {
        "type": "string"
    "bootparams.vlan": {
        "type": "string"
    "cores.dos": {
        "type": "string"
    },
    "cores.forwarding": {
        "type": "string"
    "cores.transcoding": {
```

```
"type": "string"
} ,
"cores.useSibling": {
   "type": "string"
"entitlements.FEATURE ACCOUNTING": {
   "type": "string"
"entitlements.FEATURE ADMIN SECURITY": {
   "type": "string"
"entitlements.FEATURE BFD": {
    "type": "string"
"entitlements.FEATURE IPV6": {
   "type": "string"
"entitlements.FEATURE_IWF": {
    "type": "string"
"entitlements.FEATURE QOS": {
   "type": "string"
"entitlements.FEATURE R226": {
   "type": "string"
"entitlements.FEATURE SAG": {
    "type": "string"
"entitlements.FEATURE SESSION RECORDING": {
   "type": "string"
"entitlements.Policy Server": {
    "type": "string"
} ,
"entitlements.Product Type": {
   "type": "string"
"entitlements.Routing": {
   "type": "string"
"entitlements.capacity": {
    "type": "string"
},
"entitlements.xcode-amr-cap": {
   "type": "string"
"entitlements.xcode-amrwb-cap": {
    "type": "string"
"entitlements.xcode-evs-cap": {
   "type": "string"
"entitlements.xcode-opus-cap": {
   "type": "string"
},
```

```
"entitlements.xcode-silk-cap": {
           "type": "string"
        "format.disk": {
            "type": "string"
        "license.key1": {
            "type": "string"
        "mactab.s0p0": {
           "type": "string"
        "mactab.s0p1": {
           "type": "string"
        "mactab.s1p0": {
            "type": "string"
        "mactab.s1p1": {
            "type": "string"
        },
        "mactab.wancom0": {
           "type": "string"
        "mactab.wancom1": {
            "type": "string"
        "mactab.wancom2": {
           "type": "string"
        "partitions.app": {
            "type": "string"
        } ,
        "partitions.cpp": {
            "type": "string"
        "passwords.adminPass": {
           "type": "string"
        "passwords.userPass": {
            "type": "string"
    "type": "object"
}
Example of meta_data.json
    "admin pass": "QqTUH3eFqAY8",
    "availability zone": "nova",
    "devices": [],
    "hostname": "sbc-vm.novalocal",
    "launch index": 0,
```



"meta": {

```
"bootparams.console": "VGA",
        "bootparams.gateway": "10.1.1.1",
        "bootparams.hostname": "sbc-vm",
        "bootparams.ip0": "10.1.1.2",
        "bootparams.ip6": "3fff:ac4:8001::ac4:a523/64",
        "bootparams.ip6gateway": "3fff:ac4:8001::ac4:8001",
        "bootparams.netmask0": "255.255.225.0",
        "bootparams.vlan": "0",
        "cores.dos": "1",
        "cores.forwarding": "1",
        "cores.transcoding": "0",
        "cores.useSibling": "disabled",
        "entitlements.FEATURE ACCOUNTING": "enabled",
        "entitlements.FEATURE ADMIN SECURITY": "disabled",
        "entitlements.FEATURE BFD": "disabled",
        "entitlements.FEATURE IPV6": "disabled",
        "entitlements.FEATURE IWF": "disabled",
        "entitlements.FEATURE QOS": "disabled",
        "entitlements.FEATURE R226": "disabled",
        "entitlements.FEATURE SAG": "disabled",
        "entitlements.FEATURE SESSION RECORDING": "disabled",
        "entitlements.Policy Server": "disabled",
        "entitlements.Product Type": "Session Border Controller",
        "entitlements.Routing": "disabled",
        "entitlements.capacity": "10000",
        "entitlements.xcode-amr-cap": "0",
        "entitlements.xcode-amrwb-cap": "0",
        "entitlements.xcode-evs-cap": "0",
        "entitlements.xcode-opus-cap": "0",
        "entitlements.xcode-silk-cap": "0",
        "format.disk": "true",
        "license.key1":
"abcdefghijklmnopgrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ0123456789abcdefghijklmno
pqrstuvwxyz0123456789ABCDE",
        "license.kev2":
"abcdefghijklmnopgrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ0123456789abcdefghijklmno
pgrstuvwxyz0123456789ABCDE",
        "mactab.s0p0": "fa:16:3f:15:8f:30",
        "mactab.s0p1": "none",
        "mactab.s1p0": "fa:16:3f:11:ee:4c",
        "mactab.s1p1": "none",
        "mactab.wancom0": "fa:16:3f:2d:13:94",
        "mactab.wancom1": "none",
        "mactab.wancom2": "none",
        "partitions.app": "50%",
        "partitions.cpp": "40%",
        "passwords.adminPass": "b44dda1dadd351948fcace1856ed97366e679239",
        "passwords.userPass": "b44dda1dadd351948fcace1856ed97366e679239"
    "name": "sbc-vm",
    "network config": {
        "content path": "/content/0000",
        "name": "network config"
    },
    "project id": "cf6ffe0ffaa1430497597d8756d61cf0",
    "random seed":
```

```
"0YSCvaLjZDKp2rNV67ZBI3iMuYaBiULOGQKndRlzXN3IkI8XIctZo4U5iKD81g2auJCBGeAJSeEpA +RgfUcqTY+NvTbQdX8rQDV3meB6kOr0QAQgotGk9Xtnix8rAhQW2iEUkJXG1mOl0MUzGp5Pfg40gwbqlwKjj3t1WMpgoCFfODUAu4d/n2MSKjXuJh1Xz6eSdBMJsvJ6QQlOXPut/VPNXx628=", "uuid": "71a99c78-9518-4581-88b2-15c6bd6e19f7"
}
```

Metadata Parameters

Table B-1 Supported Metadata Objects

Description		
bootparams.hostname The VM's hostname. For HA, the primary and secondary hostnames must be unique. bootparams.netmask0 The IPv4 netmask for wancom0. Derived from the Netmask of the OpenStack network connected to the wancom0 interface. bootparams.gateway The IPv4 default gateway for wancom0. bootparams.ip0 The IPv4 address for wancom0. bootparams.vlans Set this value only if the VM needs to aware of the VLAN ID for wancom0. bootparams.ip6 The IPv6 address with prefix for wancom0. Derived from the Newtron Port allocated for wancom0. Use CIDR notation, for example, 2001:DB8::1/48. bootparams.ip6gateway The IPv6 gateway address. Leave blank if not using IPv6. mactab.wancom0 The MAC address for each interface. Derived from the network port allocated for each interface. Required to map vNICs to their corresponding roles as management and media interfaces. Required to map vNICs to their corresponding roles as management and media interfaces. BHA1 hashes for the user and admin accounts. Allows you to configure the product. For example: "entitlements.Product Type": "Session Border Controller" entitlements.capacity": "10000" Ignored if Product Type is not set. Number of cores to assign to DOS, Forwarding, and Transcoding.	Metadata Key	Description
For HA, the primary and secondary hostnames must be unique. bootparams.netmask0 The IPv4 netmask for wancom0. Derived from the Netmask of the OpenStack network connected to the wancom0 interface. bootparams.gateway The IPv4 address for wancom0. bootparams.ulans Set this value only if the VM needs to aware of the VLAN ID for wancom0. bootparams.ip6 The IPv6 address with prefix for wancom0. Derived from the Newtron Port allocated for wancom0. Use CIDR notation, for example, 2001: DB8::1/48. bootparams.ip6gateway The IPv6 gateway address. Leave blank if not using IPv6. mactab.wancom0 The MAC address for each interface. Derived from the network port allocated for each interface. Required to map vNICs to their corresponding roles as management and media interfaces. Required to map vNICs to their corresponding roles as management and media interfaces. SHA1 hashes for the user and admin accounts. Passwords.userPass passwords.adminPass entitlements.Product Type Allows you to configure the product. For example: "entitlements.Product Type":"Session Border Controller" Allows you to set the entitlements for this instance. For example: "entitlements.capacity": "10000" Ignored if Product Type is not set. Number of cores to assign to DOS, Forwarding, and Transcoding.	bootparams.console	The console to which the bootparams print.
Netmask of the OpenStack network connected to the wancom0 interface. bootparams.gateway The IPv4 default gateway for wancom0. bootparams.vlans Set this value only if the VM needs to aware of the VLAN ID for wancom0. bootparams.vlans Set this value only if the VM needs to aware of the VLAN ID for wancom0. bootparams.ip6 The IPv4 address with prefix for wancom0. Derived from the Newtron Port allocated for wancom0. Use CIDR notation, for example, 2001: DB8::1/48. bootparams.ip6gateway The IPv6 gateway address. Leave blank if not using IPv6. mactab.wancom0 The MAC address for each interface. Derived from the network port allocated for each interface. Required to map vNICs to their corresponding roles as management and media interfaces. Required to map vNICs to their corresponding roles as management and media interfaces. SHA1 hashes for the user and admin accounts. PALIOWS you to configure the product. For example: "entitlements.Product Type": "Session Border Controller" Allows you to set the entitlements for this instance. For example: "entitlements.capacity": "10000" Ignored if Product Type is not set. Number of cores to assign to DOS, Forwarding, and Transcoding.	bootparams.hostname	For HA, the primary and secondary hostnames
bootparams.ip0 Dootparams.vlans Set this value only if the VM needs to aware of the VLAN ID for wancom0. Set this value only if the VM needs to aware of the VLAN ID for wancom0. The IPv6 address with prefix for wancom0. Derived from the Newtron Port allocated for wancom0. Use CIDR notation, for example, 2001:DB8::1/48. Dootparams.ip6gateway The IPv6 gateway address. Leave blank if not using IPv6. The MAC address for each interface. Derived from the network port allocated for each interface. Required to map vNICs to their corresponding roles as management and media interfaces. Required to map vNICs to their corresponding roles as management and media interfaces. SHA1 hashes for the user and admin accounts. Allows you to configure the product. For example: "entitlements.Product Type": "Session Border Controller" Allows you to set the entitlements for this instance. For example: "entitlements.capacity": "10000" Ignored if Product Type is not set. Number of cores to assign to DOS, Forwarding, and Transcoding.	bootparams.netmask0	Netmask of the OpenStack network connected to
bootparams.vlans Set this value only if the VM needs to aware of the VLAN ID for wancom0. bootparams.ip6 The IPv6 address with prefix for wancom0. Derived from the Newtron Port allocated for wancom0. Use CIDR notation, for example, 2001:DB8::1/48. bootparams.ip6gateway The IPv6 gateway address. Leave blank if not using IPv6. mactab.wancom0 mactab.wancom1 mactab.wancom2 mactab.s0p0 mactab.s1p0 mactab.s1p1 passwords.userPass passwords.userPass entitlements.Product Type Allows you to configure the product. For example: "entitlements.Product Type":"Session Border Controller" entitlements.centitlement-name> Allows you to set the entitlements for this instance. For example: "entitlements.capacity": "10000" Ignored if Product Type is not set. Number of cores to assign to DOS, Forwarding, and Transcoding.	bootparams.gateway	The IPv4 default gateway for wancom0.
bootparams.ip6 The IPv6 address with prefix for wancom0. Derived from the Newtron Port allocated for wancom0. Use CIDR notation, for example, 2001:DB8::1/48. bootparams.ip6gateway The IPv6 gateway address. Leave blank if not using IPv6. mactab.wancom0 mactab.wancom1 mactab.wancom2 mactab.s0p0 mactab.s1p0 mactab.s1p1 passwords.userPass passwords.userPass entitlements.Product Type Allows you to configure the product. For example: "entitlements.Product Type":"Session Border Controller" entitlements.centitlement-name> Allows you to set the entitlements for this instance. For example: "entitlements.capacity": "10000" Ignored if Product Type is not set.	bootparams.ip0	The IPv4 address for wancom0.
from the Newtron Port allocated for wancom0. Use CIDR notation, for example, 2001:DB8::1/48. bootparams.ip6gateway The IPv6 gateway address. Leave blank if not using IPv6. mactab.wancom0 mactab.wancom1 mactab.wancom2 mactab.s0p0 mactab.s1p0 mactab.s1p0 mactab.s1p1 passwords.userPass passwords.adminPass entitlements.Product Type Allows you to configure the product. For example: "entitlements.Product Type":"Session Border Controller" Allows you to set the entitlements for this instance. For example: "entitlements.capacity": "10000" Ignored if Product Type is not set. Number of cores to assign to DOS, Forwarding, and Transcoding.	bootparams.vlans	
mactab.wancom0 mactab.wancom1 mactab.wancom2 mactab.sop0 mactab.sop0 mactab.s1p0 mactab.s1p1 passwords.userPass passwords.adminPass entitlements.Product Type allows you to configure the product. For example: "entitlements.Product Type": "Session Border Controller" Allows you to set the entitlements for this instance. For example: "entitlements.capacity": "10000" Ignored if Product Type is not set. Number of cores to assign to DOS, Forwarding, and Transcoding.	bootparams.ip6	from the Newtron Port allocated for wancom0. Use
mactab.wancom1 mactab.wancom2 mactab.s0p0 mactab.s1p0 mactab.s1p1 passwords.userPass passwords.adminPass entitlements.Product Type entitlements. <entitlement-name> Allows you to set the entitlements for this instance. For example: "entitlements.capacity": "10000" Ignored if Product Type is not set. Cores.dos cores.forwarding the network port allocated for each interface. Required to map vNICs to their corresponding roles as management and media interfaces. Required to map vNICs to their corresponding roles as management and media interfaces. Required to map vNICs to their corresponding roles as management and media interfaces. Required to map vNICs to their corresponding roles as management and media interfaces. Required to map vNICs to their corresponding roles as management and media interfaces. SHA1 hashes for the user and admin accounts. Passwords.adminPass Allows you to configure the product. For example: "entitlements.Product Type": "Session Border Controller" Ignored if Product Type is not set. Cores.dos cores.forwarding</entitlement-name>	bootparams.ip6gateway	
mactab.wancom2 mactab.s0p0 mactab.s1p0 mactab.s1p0 mactab.s1p1 passwords.userPass passwords.adminPass entitlements.Product Type allows you to configure the product. For example: "entitlements.Product Type":"Session Border Controller" Allows you to set the entitlements for this instance. For example: "entitlements.capacity": "10000" Ignored if Product Type is not set. cores.dos cores.forwarding Required to map vNICs to their corresponding roles as management and media interfaces. Required to map vNICs to their corresponding roles as management and media interfaces. SHA1 hashes for the user and admin accounts. Product. For example: "entitlements.Product Type": "Session Border Controller" Ignored if Product Type is not set. Number of cores to assign to DOS, Forwarding, and Transcoding.	mactab.wancom0	The MAC address for each interface. Derived from
mactab.s0p0 mactab.s1p0 mactab.s1p1 passwords.userPass passwords.adminPass entitlements.Product Type Allows you to configure the product. For example: "entitlements.Product Type":"Session Border Controller" Allows you to set the entitlements for this instance. For example: "entitlements.capacity": "10000" Ignored if Product Type is not set. cores.dos cores.forwarding	mactab.wancom1	1
mactab.supu mactab.s1p0 mactab.s1p1 passwords.userPass passwords.adminPass entitlements.Product Type Allows you to configure the product. For example: "entitlements.Product Type":"Session Border Controller" Allows you to set the entitlements for this instance. For example: "entitlements.capacity": "10000" Ignored if Product Type is not set. cores.dos cores.forwarding and Transcoding.	mactab.wancom2	
mactab.s0p1 mactab.s1p1 passwords.userPass passwords.adminPass entitlements.Product Type Allows you to configure the product. For example: "entitlements.Product Type":"Session Border Controller" Allows you to set the entitlements for this instance. For example: "entitlements.capacity": "10000" Ignored if Product Type is not set. cores.dos cores.forwarding Allows you to set the ontitlements for this instance. For example: "entitlements.capacity": "10000"	mactab.s0p0	roles as management and media interfaces.
passwords.userPass passwords.adminPass entitlements.Product Type Allows you to configure the product. For example: "entitlements.Product Type":"Session Border Controller" Allows you to set the entitlements for this instance. For example: "entitlements.capacity": "10000" Ignored if Product Type is not set. Cores.dos cores.forwarding Number of cores to assign to DOS, Forwarding, and Transcoding.	mactab.s1p0	
passwords.userPass passwords.adminPass entitlements.Product Type Allows you to configure the product. For example: "entitlements.Product Type":"Session Border Controller" Allows you to set the entitlements for this instance. For example: "entitlements.capacity": "10000" Ignored if Product Type is not set. cores.dos cores.forwarding Number of cores to assign to DOS, Forwarding, and Transcoding.	mactab.s0p1	
entitlements.Product Type Allows you to configure the product. For example: "entitlements.Product Type":"Session Border Controller" Allows you to set the entitlements for this instance. For example: "entitlements.capacity": "10000" Ignored if Product Type is not set. cores.dos cores.forwarding Number of cores to assign to DOS, Forwarding, and Transcoding.	mactab.s1p1	
entitlements.Product Type Allows you to configure the product. For example: "entitlements.Product Type":"Session Border Controller" Allows you to set the entitlements for this instance. For example: "entitlements.capacity": "10000" Ignored if Product Type is not set. Cores.dos Cores.forwarding Number of cores to assign to DOS, Forwarding, and Transcoding.	passwords.userPass	SHA1 hashes for the user and admin accounts.
"entitlements.Product Type":"Session Border Controller" Allows you to set the entitlements for this instance. For example: "entitlements.capacity": "10000" Ignored if Product Type is not set. Cores.dos Cores.forwarding Number of cores to assign to DOS, Forwarding, and Transcoding.	passwords.adminPass	
entitlements. <entitlement-name> Allows you to set the entitlements for this instance. For example: "entitlements.capacity": "10000" Ignored if Product Type is not set. Cores.dos Cores.forwarding Number of cores to assign to DOS, Forwarding, and Transcoding.</entitlement-name>	entitlements.Product Type	Allows you to configure the product. For example:
For example: "entitlements.capacity": "10000" Ignored if Product Type is not set. cores.dos cores.forwarding Number of cores to assign to DOS, Forwarding, and Transcoding.		_ ==
Ignored if Product Type is not set. cores.dos	entitlements. <entitlement-name></entitlement-name>	
cores.dos Number of cores to assign to DOS, Forwarding, and Transcoding.		"entitlements.capacity": "10000"
cores.dos Number of cores to assign to DOS, Forwarding, and Transcoding.		Ignored if Product Type is not set.
cores.forwarding and Transcoding.	cores.dos	 -
cores.transcoding	cores.forwarding	
	cores.transcoding	



Table B-1 (Cont.) Supported Metadata Objects

Metadata Key	Description
cores.useSibling	If SMT topology is exposed by the hypervisor, you can enable sibling cores for datapath usage. Disabled by default.
license.key1	The set of license keys.
license.key2	
license.key <n></n>	
partitions. <name1> partitions.<name2> partitions.<name3></name3></name2></name1>	Data disk partitions to be created during deployment. The maximum number of partitions are 4, and the value should be the percentage of the disk to be allocated. For example:
partitions. <name4></name4>	
	"partitions.app": "20%",
	"partitions.cdr": "80%",
format.disk	Set to "true" to enable formatting the SBC's system-disk and data-disk. Parameter is valid only if the partitions have been specified. Set to "false" to format only the data-disk.

User Data Tokens

Table B-2 User Data Tokens

S. No	Token Name	Token value
1	\${hostname}	Target hostname of the VM
2	\${forwardingCores}	Number of forwarding cores
3	\${dosCores}	Number of DoS cores
4	\${transcodingCores}	Number of transcoding cores
5	\${useSiblingCoreDatapath}	Enable or disable sibling cores for datapath usage
6	\${snmpCommunityName}	SNMP community name
7	\${snmplpAddress}	SNMP IP address in dot-decimal notation
8	\${ntpServer1}	First NTP IP address in dot-decimal notation
9	\${ntpServer2}	Second NTP IP address in dot- decimal notation

Example Configuration

The example configuration shows the tokens that will be replaced during the deployment of a virtual SBC. Tokens use the syntax $\{token-name\}$.

```
<config lastObjectId='11' schemaVersion='ECZ9.0.0/0' versionStr='ECZ9.0.0 GA
(Build 59)'>
    <systemConfig hostname='${hostname}'
    descr=''
    location=''</pre>
```



```
sysContact=''
 sysName=''
 sysLocation=''
 acp-tls-profile=''
 enableSnmp='enabled'
 snmpAuth='disabled'
 slogNotification='disabled'
 enableSnmpMonitor='disabled'
 enableEnvMonitor='disabled'
 enableMblkTracking='disabled'
 enableL2MissReport='enabled'
 slogHistLen='1'
 slogLevel='WARNING'
 systemLogLevel='WARNING'
 processLogLevel='NOTICE'
 logIpAddr='0.0.0.0'
 logPort='0'
 callTrace='disabled'
 internalTrace='disabled'
 logFilter='all'
 defaultGW=''
 reboot='enabled'
 reboot-exceptions=''
 telnet-timeout='0'
 console-timeout='0'
 remote-control='enabled'
 cliAudit='enabled'
 sourcerouting='disabled'
 cliMore='disabled'
 cliMoreHeight='24'
 debug-timeout='0'
 trapEventLifeTime='0'
 idsSyslogFacility='-1'
 options=''
 defaultGWv6=''
 ipv6Mtu='1500'
 ipv4Mtu='1500'
 ipv6Support='enabled'
 cleanupTimeOfDay='00:00'
 snmpEngineIDSuffix=''
 snmpAgentMode='v1v2'
 dpdkForwardCores='${forwardingCores}'
 dpdkDosCores='${dosCores}'
 dpdkXcodeCores='${transcodingCores}'
 siblingCoreDpdkUsage='${useSiblingCoreDatapath}'
 lastModifiedBy='admin@console'
 lastModifiedDate='2018-01-16 17:46:11' objectId='1'>
<collectConfig sampleInt='5'</pre>
 pushInt='15'
 mode='disabled'
 startTime='0'
 endTime='0'
 redCollectPort='0'
 redNumTrans='1000'
 redSyncStartTime='5000'
 redSyncCompTime='1000'
```

```
pushSuccessTrapEnabledState='disabled'/>
 <commMonitor state='disabled'</pre>
    sbcGrpId='0'
    tlsProfile=''
    enableQos='enabled'
    interimQoSUpdate='disabled'/>
 </systemConfig>
 <snmpCommunity communityName='${snmpCommunityName}'</pre>
    accessMode='READ-ONLY'
    lastModifiedBy='admin@10.196.0.194'
    lastModifiedDate='2018-01-24 22:40:59' objectId='7'>
    <key>${snmpCommunityName}</key>
    <ipAddresses addr='${snmpIpAddress}' numBits='0'/>
 </snmpCommunity>
 <ntpConfig lastModifiedBy='admin@10.196.0.18'</pre>
    lastModifiedDate='2018-03-28 16:55:15' objectId='11'>
    <servers name='${ntpServer1}'/>
    <servers name='${ntpServer2}'/>
 </ntpConfig>
</config>
```

