Oracle® Communications Application Session Controller

System Installation and Commissioning Guide

Release 3.7.0M4

July 2016

ORACLE

Copyright ©2016, 2005, Oracle and/or its affiliates. All rights reserved.

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

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

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

U.S. GOVERNMENT END USERS: Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

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

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

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

This software or hardware and documentation may provide access to or information on 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. 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.

Contents

Preface 7
About OS-E® Documentation 7
Revision History 8
Conventions Used in This Manual 9
Typographical Conventions 9
Acronyms 9

OS-E® Series Overview 15
About this Chapter 15
Running the OS-E on Other Devices 15
Information on OS-E Software and Licensing 16
Obtaining Your License 16
License Expirations and Renewals 17
System Management 17

Installing And Upgrading The OS-E System 19
About This Chapter 19
Installing the OS-E 19
Installing Oracle Linux 7 20
Obtaining the OS-E Installation File 27
Configuring a Yum Repository 30
Installing the OS-E Appliance 31
Upgrading To 3.7.0M4 From Release 3.7.0M3 32
Backing Up the OS-E Configuration, Files, and Databases 32

Installing Oracle Linux 7 33
Installing The OS-E 33
Restoring Your Configuration, Files, and Databases On The OS-E 33

Quick Commissioning New OS-E Systems 35

About This Chapter 35
Prerequisites to Quick Commissioning 35
Building the Configuration File 36
Basic Network Topology 37
Step 1. Configuring Basic IP Connectivity 38
Using the Setup Script 39
Enabling Network Access 40
Defining a Default Route and Gateway IP 40
Launching the OS-E Management System 40
Changing the Linux Root Password 41
Step 2. Configuring Advanced IP Connectivity 42
Step 3. Creating User Accounts for Basic Access 44
Step 4. Enabling Master Services 46
Step 5. Configuring Basic Services 47
Step 6. Enabling the Virtual System Partition (VSP) 49
Step 7. Configuring the Accounting Environments 50
Step 8. Editing the Default Session Configuration 52
Step 9. Enabling Registration Services 53
Step 10. Reviewing the Configuration 54

Installing OS-E Clusters 57

About This Chapter 57
OS-E Cluster Overview 57
Cluster Operations and Services 58
Master-Services 58
Heartbeat Interface, BOOTP, and Messaging 63
Event Logging 63
Network Time Protocol (NTP) 64

Cluster Redundancy Operations 64
Notes on Cluster Management 64
Cluster Installation Prerequisites 65
Cluster Installation Procedure 66
Configuring External Messaging 72
Configuring Cluster Load Balancing 73
Restarting an OS-E Cluster 74

Installing Certificates and Commissioning TLS Networks 77

About This Chapter 77
TLS Overview 77
Steps to Configuring TLS 77
Before Configuring TLS 79
Step 1. Creating a Self-Signed Certificate and Key Pair from the OS-E 80
Step 2. Generating a Certification Signing Request 84
Step 3. Signing a CSR Using Either a Valid CA or OpenSSL 85
Step 4. Updating the Self-Signed Certificate 97
Subject Alternative Name for HTTPS Certificates Support 97
Configuring the Certificate on the OS-E 100
Displaying the Certificates Installed on the OS-E 100
Other TLS Certificate Settings 100
Using Certificate vs. Default-Outgoing-Settings 100
Verifying Peer Certificates 101
Enabling Peer Certificate Verification 102
Controlling the CA Files and CRLs to Apply to the Certificate 103
Setting the Required Peer Name 103
Configuring TLS on Ethernet Interfaces 103

Configuring Secure Media (SRTP) Sessions 105

About This Chapter 105
Anchoring Media Sessions 105
Configuring Inbound and Outbound Encryption 106
Inbound Encryption Mode and Type 107

Outbound Encryption Mode, Type, and Require-TLS Setting 108
Linksys Encryption 109

Creating and Commissioning USB Sticks 113

Supported USB Sticks 113
USB Stick Restrictions 113
Important Note About the New USB Stick 114
Creating a New USB Rescue Stick 115
Using the Rescue Utility USB 115
Using the Rescue Mode 118
System and Data Drive Locations 119

Installing and Running the OS-E Virtual Machine 121

Server-Based Requirements 121
Linux Installations 122
Installing the VM 122
Installing the OS-E on an Oracle Virtual Machine 122
Installing the OS-E On a VMware ESXi 143
Installing the OS-E As a XEN Virtual Machine 146
Installing the OS-E On KVM 151
Configuring the VM 151
Using Config Setup 152
Sample VM Configuration 153
Enabling the OS-E Management System 159
Bridging to Additional Ethernet Ports 159
Adding an Additional VMnet 159
Editing the VM Configuration File 160
OS-E-VM Troubleshooting 161
Installing the VM on Slow Systems 161
Other VM Limitations and Considerations 164

Preface

About OS-E® Documentation

The OS-E references in this documentation apply to the OS-E operating system software that is used for the following Oracle and third-party SBC products:

- Oracle Communications Application Session Controller (ASC)
- Oracle Communications WebRTC Session Controller Media Engine (WSC-ME)
- Oracle Communications Covergence Access Controller (OCCAS) Media Engine (OCCAS-ME)
- Oracle Communications OS-E Session Director (SD) Session Border Controller (SBC)
- Oracle Communications 2600 Session Director (SD) Session Border Controller (SBC)
- Third-party products that license and use Oracle Communications OS-E software on an OEM basis.

Unless otherwise stated, references to OS-E in this document apply to all of the Oracle and third-party vendor products that use OS-E software.

The following documentation set supports the current release of the OS-E software.

- *Oracle Communications Application Session Controller System and Installation Commissioning Guide*
- *Oracle Communications Application Session Controller Management Tools*
- *Oracle Communications Application Session Controller System Administration Guide*
- *Oracle Communications Application Session Controller Session Services Configuration Guide*

- *Oracle Communications Application Session Controller Objects and Properties Reference*
- *Oracle Communications Application Session Controller System Operations and Troubleshooting*
- *Oracle Communications Application Session Controller Release Notes*
- *Oracle Communications Application Session Controller Single Number Reach Application Guide*
- *Oracle Communications Application Session Controller Web Services SOAP REST API*
- *Oracle Communications WebRTC Session Controller Installation Guide*

Revision History

This section contains a revision history for this document.

Date	Revision Number	Description
May, 2016	Revision 1.00	• GA Release of OS-E 3.7.0M4 software.
July, 2016	Revision 1.01	• Updates the “Configuring a Yum Repository” procedure. • Updates the version number of Oracle Linux on which you can install the OS-E.

Conventions Used in This Manual

Typographical Conventions

Key Convention	Function	Example
key name	Identifies the name of a key to press.	Type **abc**, then press [ENTER]
CTRL+*x*	Indicates a control key combination.	Press CTRL+C
brackets [ ]	Indicates an optional argument.	[*portNumber*]
braces { }	Indicates a required argument with a choice of values; choose one.	{enabled | disabled}
vertical bar |	Separates parameter values. Same as "or."	{TCP | TLS}
Monospaced bold	In screen displays, indicates user input.	config> **`config vsp`**
Monospaced italic	In screen displays, indicates a variable—generic text for which you supply a value.	`config servers>` **`config lcs`** *`name`*
bold	In text, indicates literal names of commands, actions, objects, or properties.	...set as the secondary directory service (with the **unifier** property)...
bold italic	In text, indicates a variable.	...set the **domain** property of the ***directory*** object.

Acronyms

The OS-E manuals contain the following industry-standard and product-specific acronyms:

AAA	Authentication, authorization, and accounting
ALI	Automatic location identifier
ANI	Automatic number identification
ANSI	American National Standards Institute
AOR	Address of record
API	Application programming interface
ARP	Address Resolution Protocol
AVERT	Anti-virus emergency response team

B2BUA	Back-to-back user agent
BOOTP	Bootstrap Protocol
CA	Certificate authority
CAP	Client application protocol
CBC	Cipher block chaining
CBN	Call back number
CCS	Converged Communication Server
CDR	Call detail record
CIDR	Classless interdomain routing
CLI	Command line interface
CMOS	Comparison mean opinion score
CNAME	Canonical name record
CNI	Calling number identification
CODEC	Compressor/decompressor or coder/decoder
CPE	Customer-premise equipment
CRL	Certificate revocation list
CSR	Certificate signing request
CSTA	Computer-supported telecommunications applications
CSV	Comma-separated values
DDDS	Dynamic delegation discovery system
DHCP	Dynamic Host Configuration Protocol
DMZ	Demilitarized zone
DN	Distinguished name
DNIS	Dialed number identification service
DNS	Domain name service
DOS	Denial of service
EIM	Enterprise instant messaging
ESD	Electrostatic discharge
ESGW	Emergency services gateway
ESQK	Emergency services query key
ESRN	Emergency services routing number
FQDN	Fully qualified domain name

GUI	Graphical user interface
HTTP	Hypertext Transfer Protocol
HTTPS	Hypertext Transfer Protocol Secure
I2	National Emergency Number Association defined VoIP solution
ICAP	Internet Calendar Access Protocol
ICMP	Internet Control Message Protocol
IM	Instant messaging
IP	Internet Protocol
JDBC	Java database connectivity
JMX	Java management extensions
JRE	Java runtime environment
LATA	Local access and transport area
LCS	Live Communications Server
LCR	Least-cost routing
LDAP	Lightweight Directory Access Protocol
LIS	Location information service
MAC	Media access control
MCS	Multimedia Communications Server
MIB	Management information base
MOS	Mean opinion score
MSAG	Master street address guide
MTU	Maximum transmission unit
NAPTR	Naming authority pointer
NAT	Network address translation
NENA	National Emergency Number Association
NIC	Network interface card
NS	Name server
NSE	Named signaling events
NTLM	NT Lan Manager
NTP	Network Time Protocol
OC	Office Communicator
OCI	Open Client Interface

ODBC	Open database connectivity
OTP	Over temperature protection
OVP	Over voltage protection
PBX	Private branch eXchange
PEM	Privacy-enhanced mail
PERL	Practical Extraction and Reporting Language
PING	Packet internet groper
PKCS#12	Public Key Cryptography Standard #12
PKI	Public Key Infrastructure
PSAP	Public safety answering point
PSCP	PuTTY secure copy
PSTN	Public switched telephone network
QOP	Quality of protection
QOS	Quality of service
RADIUS	Remote Authentication Dial-in User Service
RTC	Real-time collaboration
RTCP	Real-time Control Protocol
RTP	Real-time Transport Protocol
RTT	Round-trip time
SATA	Serial ATA
SCSI	Small computer system interface
SDK	Software development kit
SDP	Session Description Protocol
SFTP	Secure Shell File Transfer Protocol
SIMPLE	SIP Instant Messaging and Presence Leveraging Extension
SIP	Session Initiation Protocol
SIPS	Session Initiation Protocol over TLS
SLB	Server load balancing
SMB	Server message block
SNMP	Simple Network Management Protocol
SOA	Server of authority
SOAP	Simple Object Access Protocol

SQL	Structured Query Language
SRTP	Secure Real-time Transport Protocol
SRV	Server resource
SSH	Secure Shell
SSL	Secure socket layer
SSRC	Synchronization source
STUN	Simple Traversal of UDP over NATs
TCP	Transmission Control Protocol
TDM	Time division multiplexing
TGRP	Trunk group
TLS	Transport Layer Security
TOS	Type of service
TTL	Time to live
UPS	Uninterruptable power supply
US	User agent
UAC	User agent client
UAS	User agent server
UDP	User Datagram Protocol
UID	Unique identifier
URI	Uniform resource identifier
URL	Uniform resource locator
UTC	Universal coordinated time
VoIP	Voice over IP
VLAN	Virtual local area network
VPC	VoIP positioning center
VRRP	Virtual Router Redundancy Protocol
VSP	Virtual system partition
VXID	Virtual router interface ID
WAR	Web application resource
WAV	Waveform audio
WM	Windows Messenger
WSDL	Web Services Description Language

XML	Extensible Markup Language
XSL	Extensible Stylesheet Language

Chapter 1. OS-E® Series Overview

About this Chapter

This chapter provides an overview of the Oracle OS-E® Series software. In addition to running on Oracle hardware offerings, OS-E software is supported on compatible third-party platforms and blades.

Running the OS-E on Other Devices

You can run the OS-E on a number of supported third-party servers.

The following platforms have been certified for use with the OS-E:

- Sun Netra X5-2
- Sun Server X5-2
- Sun Netra X3-2
- HPDL160 G9
- NN2610
- NN2620

The following VM platforms have been certified for use with the OS-E:

- OVM 3.2.8
- VMware ESXi 5.5
- XEN 3.4.3
- KVM on OL7

Information on OS-E Software and Licensing

Using the Internet and secure Web URLs, Oracle provides all necessary software downloads for USB creation, product licensing, and commissioning of your selected hardware.

As part of each download, and depending on your actual requirements, Oracle can provide the following:

- Oracle USB Creation Utility with the OS-E software
- Documentation on how to create an Oracle USB stick and commission the OS-E software on your selected hardware
- Standard set of Oracle OS-E technical publications
- Royalty-bearing codec licenses (if purchased)

You must provide a USB stick with 4GB storage to handle Oracle software downloads. Oracle has tested a variety of USB sticks available from current suppliers and manufacturers. Most USB sticks manufactured today will work.

For complete information on accessing the Oracle download server, creating an installation USB stick, and commissioning OS-E systems, refer to Chapter 7, Creating and Commissioning USB Sticks.

Obtaining Your License

If you are running OS-E version 3.7.0M2 or later and are NOT using Royalty-bearing codecs, you should begin using the default shipping license. There is no loss of functionality as a result.

The default license enables the maximum number of sessions for the system. In the past, the software stopped allowing new sessions at your specific licensed maximum depending on the license. The system no longer relies on the license to apply an upper limit. However, depending on the server hardware in use, the system may not be capable of supporting a higher number of sessions. You may want to edit your configuration file and add the parameters for the maximum number of media sessiosn to be sure that you do not exceed the capabilities of your hardware. This applies to most deployments running a small number of sessions on smaller third-party hardware that could potentially have a problem if traffic increases to a number larger than the system can handle.

Note: You may continue using the legacy licensing system or you may transition to self-provisioning entitlements. In both cases, ensure that your system's functionality abides by your organization's contractual obligations with Oracle.

The following are the royalty-bearing codecs supported by the OS-E:

- AMRWB
- AMRNB
- G723
- G729

License Expirations and Renewals

If your customer-specific license comes with an expiration date, the OS-E system will generate an event when the license nears the expiration date. Contact your Oracle Sales Representative to complete the purchase of the features that you are testing. These expiring licenses should apply only to customers testing royalty-bearing codecs in the transcoding feature of the OS-E.

System Management

Before you install the system, you should decide on the management tool(s) that you want to use to configure and monitor the system. This will help you decide where you need to create connections based on your equipment and network resources.

System management capabilities include the following secure management interfaces:

- The OS-E command line interface (CLI) from a local console, Telnet, or SSH connection
- The OS-E Management System, a graphical user interface (GUI) that supports remote management using the Internet Explorer Web browser
- Simple Network Management Protocol (SNMP) using third party SNMP MIB compiler/browser applications
- Web Services Description Language (WSDL) and Simple Object Access Protocol (SOAP) messaging using the software development kit (SDK)

For information on configuring the management options, refer to *Oracle Communications OS-E Management Tools*.

Chapter 2. Installing And Upgrading The OS-E System

About This Chapter

This chapter covers OS-E system installation and upgrading.

Installing the OS-E

Starting with release 3.7.0M4, the OS-E runs on Oracle Linux and uses yum to install and update RPM files. In prior releases, the OS-E install package came with its own custom kernel.

> **Note:** While the OS-E operates under Oracle Linux, it is not certified to operate under other Linux environments.

You must have Oracle Linux Release 7.2 or higher installed on your hardware prior to installing the OS-E.

Due to the fact that the installation and upgrade procedures are significantly different in release 3.7.0M4 than in the other 3.7.0M*x* releases, two versions of the installation guide have been created. If you are currently trying to install or upgrade a 3.7.0M*x* release prior to 3.7.0M4, see the *Oracle Communications Application Session Controller Installation and Commissioning Guide Release 3.7.0*.

> **Note:** If you are upgrading an older version of OS-E software as opposed to performing a fresh installation , Oracle recommends you create a USB rescue stick to preserve the existing configuration to install on the new software. For more information on creating a USB rescue stick, see “Creating a New USB Rescue Stick” in *Chapter 7, Creating and Commissioning USB Sticks*.

To install the OS-E you must:

- Install Oracle Linux version 7.2 or higher
- Download and copy the OS-E file to a USB stick
- Mount the OS-E software onto your hardware
- Configure a yum repository on which to point Oracle Linux
- Install the OS-E appliance

Installing Oracle Linux 7

Before you can install the OS-E, you must have Oracle Linux installed on your hardware. You can either install Oracle Linux via a USB stick or a DVD. This guide documents installing Oracle Linux via a USB.

> **Note:** When you install Oracle Linux via a USB stick, you must have a 16 GB or bigger USB drive.

For a much more comprehensive and thorough description of installing Oracle Linux 7, see *https://docs.oracle.com/cd/E52668_01/E54695/E54695.pdf*.

To install Oracle Linux via a USB stick:

1. Download “Oracle Linux 7.*X* for x86 64 bit ISO image” from *http://edelivery.oracle.com/linux*.

2. Create a bootable USB stick that contains the full Oracle Linux 7 ISO image. The following example uses Rufus 2.8 software to create the bootable USB stick.

Note: You can also install Oracle Linux via a DVD by downloading "Oracle Linux 7.1 for x86 64 bit ISO image" from *http://edelivery.oracle.com/linux* and burning the *.iso image onto a DVD.

3. Insert your bootable Oracle Linux 7 USB drive onto your hardware.

4. Boot the system from the boot image by selecting **Boot from usb** from the boot menu options.

The system locates the boot image file and the boot menu appears.

5. Select **Install Oracle Linux 7.2** and hit <*Enter*>.
6. Select the appropriate language and select **Set keyboard to default layout for selected language**. Click **Continue**.

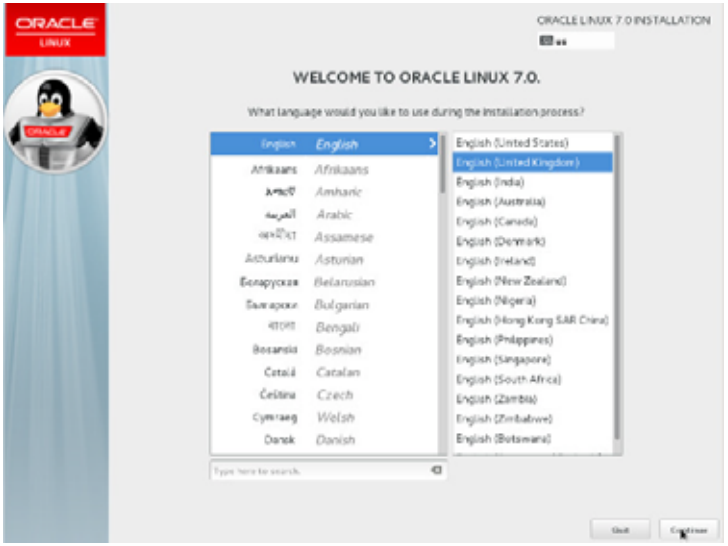

The “Installation Summary” screen appears.

7. Complete any marked items. Depending on your requirements, you may also need to alter the default settings by clicking on the relevant links.

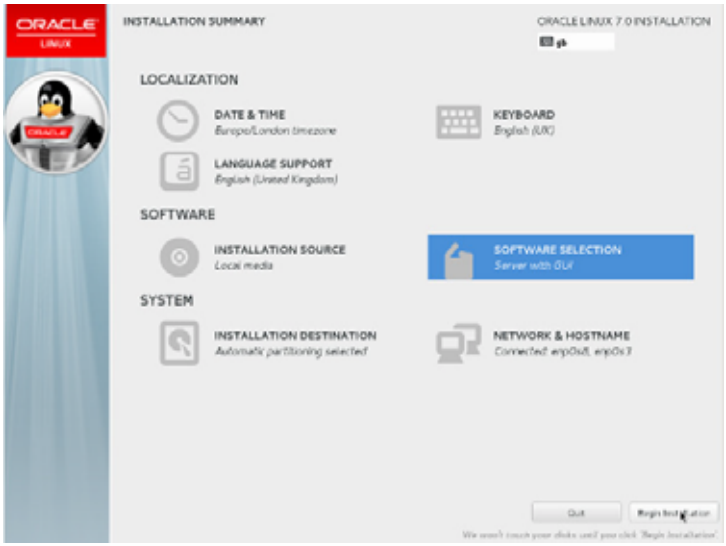

8. Click **Installation Destination**.

9. Select the local disks you want to use for the installation.

Note: The installation program does not make any changes to any of the disks and storage until you click **Begin Installation** on the “Installation Summary” screen.

- Choose the disks on which you want to install Oracle Linux from the “Local Standard Disks” section. A tick icon displays next to the selected disks.
- Ensure the **Automatically configure partitioning** option is selected (by default, this option is selected).
- At the bottom of the screen, the system displays how much disk space you need for the software you have selected. With automatic partitioning, you may not have sufficient space to install the software if the disk is already partitioned. If you need to free some disk space, select **I would like to make additional space available** and click **Done**.

  Note: You must have disk size of at least 50 G.

  The “Reclaim Disk Space” window appears.

  Note: If there is still insufficient disk space when you click **Done**, the system prompts you to free disk space.

- Once you have selected the disks you want to use, click **Delete all** to free disk space and then click **Reclaim Space**.

  For more information on configuring partitioning, see *https://docs.oracle.com/cd/E52668_01/E54695/E54695.pdf*.

10. Click **Begin Installation** once you have completed any necessary updates to the default configuration.

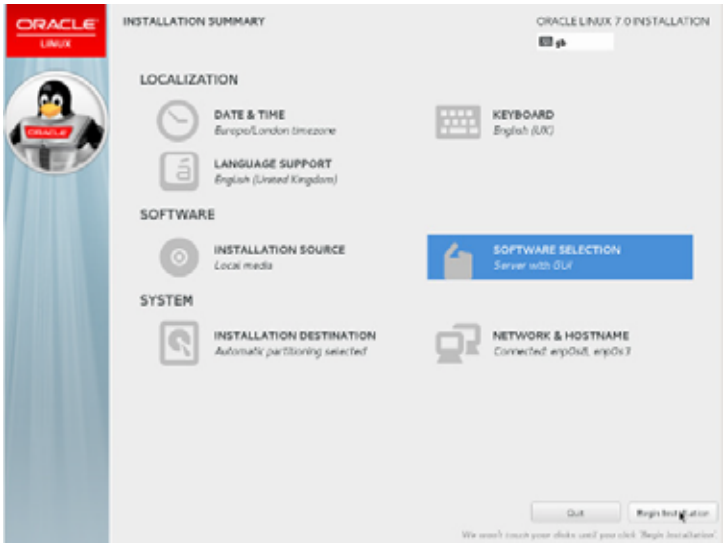

11. Click **Root Password**.

12. Enter the root password and click **Done**.

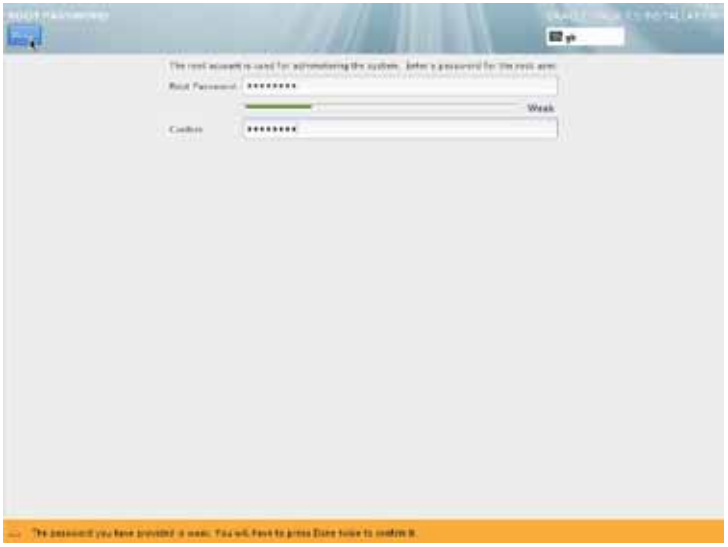

Linux installs.

13. Click **Reboot** when prompted.

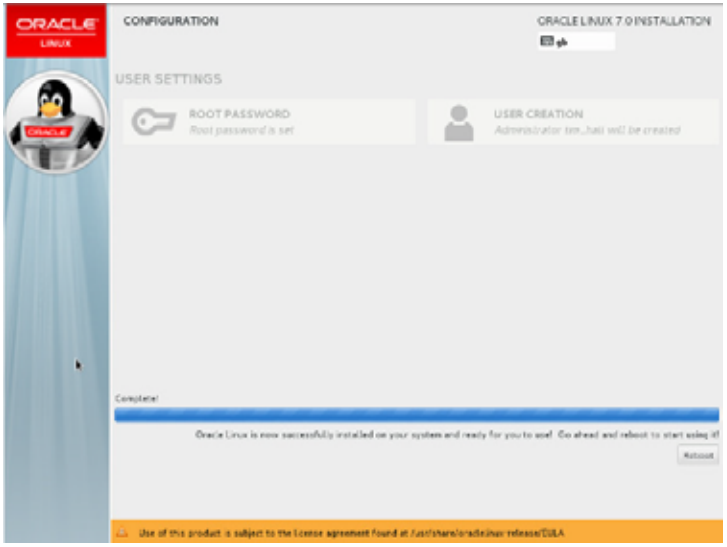

14. Run the following command to rename the Ethernet devices according to your system BIOS once you successfully complete the installation.

```
grubby --args=net.ifnames=0 --update-kernel=ALL
```

15. Reboot the system.

Obtaining the OS-E Installation File

Before you can install the OS-E, you must first download the ISO files you need and copy them to a USB stick.

Software can be downloaded from either the Oracle Software Delivery Cloud or the My Oracle Support Patches and Updates tab.

To access the Oracle Software Delivery Cloud:

1. Access the https://edelivery.oracle.com link.
2. Select the **Sign In/Register** tab and enter your **username** and **password**.

   Note: If you are a new user, you must create an account.
3. Click the checkbox to agree to the to the **Oracle Trial License Agreement and Export Restrictions** and click **Continue**.
4. Select the Oracle Communications product pack.
5. Select the Acme Packet OS platform and click **Go**.
6. Download the following file and click **Download**.
   - Oracle Communications Application Session Controller E3.7.0m4 Installation Repository
7. Copy the file onto a USB stick.

To access the Oracle Support Software Patches and Updates:

1. Log into the My Oracle Support Portal.
2. Select the **Patches and Updates** tab.

3. Select the **Search** tab and click **Product or Family (Advanced)**.

4. **Product:** Enter **Oracle Communications Application Session Controller**.
5. **Release:** Enter **Application Session Controller 3.7**.
6. Click **Search**. The available distribution formats appear and include the following information:
   - Patch Name
   - Description
   - Release
   - Platform (Language)
   - Classification
   - Product
   - Prerequisite Requirement
   - Size
   - Download Access
7. Select the distribution format that you require.
8. Click either **Download** to download the file or **Read Me** to view the Build Notes for this patch.
9. Copy the file onto a USB stick.

Mounting the OS-E Installation file.

Once you have the installation file, you must install it on your hardware.

To mount the OS-E installation file:

1. Insert the USB stick onto your hardware and locate the USB stick partition to mount (for example, /dev/sdd1).

```
[root@localhost]# sudo <fdisk -l>
```

2. Create a mount point and mount the installation file.

```
[root@localhost]# sudo mkdir  /mnt/usb
[root@localhost]# sudo mount /dev/sdd1  /mnt/usb
```

3. Extract the files via “unzip”.

   Note: The following is an example and uses example values only.

```
[root@localhost]#cd /mnt/usb
[root@localhost]# unzip 370m4p0-2016-03-24_22-51-16


Archive:  370m4p0-2016-03-24_22-51-16.zip
               inflating: 370m4p0/kernel-uekcov-debuginfo-3.8.13-
   118.4.1.370m4p0.69387.el7uekcov.x86_64.rpm
               inflating: 370m4p0/
   apache-commons-fileupload-1.3.1-4.el7.noarch.rpm
              inflating: 370m4p0/crtmpserver-690-3.el7.i686.rpm
              inflating: 370m4p0/slotmap-2.2.3-1.el7.i686.rpm
               inflating: 370m4p0/jre-8u65-fcs.1.el7.i686.rpm
              inflating: 370m4p0/
   asc-app-emblcrimport-E3.7.0.M4P0-69407.el7.i686.rpm
               inflating: 370m4p0/
   apache-commons-fileupload-javadoc-1.3.1-4.el7.noarch.rpm
              inflating: 370m4p0/portlet-2.0-api-1.0-9.el7.noarch.rpm
               inflating: 370m4p0/
   kernel-uekcov-debuginfo-common-3.8.13-118.4.1.370m4p0.69387.el7uek
   cov.x86_64.rpm
              inflating: 370m4p0/libipp-7.0-2.el7.i686.rpm
               inflating: 370m4p0/
   kernel-uekcov-covmodule-2.8.1-3.8.13.118.4.1.370m4p0.68883.el7uekc
   ov.x86_64.rpm
               inflating: 370m4p0/
   postgresql-odbc-09.03.0100-2.el7.cov1.i686.rpm
              inflating: 370m4p0/
   asc-appliance-E3.7.0.M4P0-69407.el7.i686.rpm
               inflating: 370m4p0/
   kernel-uekcov-3.8.13-118.4.1.370m4p0.69387.el7uekcov.x86_64.rpm
              inflating: 370m4p0/libunwind-1.1-5.el7.2.i686.rpm
               inflating: 370m4p0/
   kernel-uekcov-firmware-3.8.13-118.4.1.370m4p0.69387.el7uekcov.x86_
   64.rpm
              inflating: 370m4p0/
   asc-selinux-policy-E3.7.0.M4P0-29.el7.noarch.rpm
inflating: 370m4p0/asc-rescue-stick-E3.7.0.M4P0-29.el7.i686.rpm
```

```
             inflating: 370m4p0/
asc-app-samples-E3.7.0.M4P0-69407.el7.i686.rpm
             inflating: 370m4p0/
portlet-2.0-api-javadoc-1.0-9.el7.noarch.rpm
            inflating: 370m4p0/h264bitstream-0.1.6-1.el7.i686.rpm
             inflating: 370m4p0/rtmpdump-2.3-1.el7.i686.rpm
            inflating: 370m4p0/repodata/filelists.xml.gz
            inflating: 370m4p0/repodata/primary.xml.gz
            inflating: 370m4p0/repodata/repomd.xml
            inflating: 370m4p0/repodata/other.xml.gz
```

Note: If your system does not have the “unzip” package installed, execute the **yum install unzip** command.

Configuring a Yum Repository

Oracle Linux 7 uses “yum” to install and update RPM files. Yum uses a URL to point at repositories. The URL must be pointed to the mounted USB directory you create when you mount the installation file, described in Mounting the OS-E Installation file.

To configure a yum repository:

1. Create the repository entry using root permissions.

```
cat >/etc/yum.repos.d/asc.repo <<EOF
             [asc-base]
             name=Application Session Controller Base Repository
             baseurl=file:///mnt/usb/370m4p0
             gpgcheck=0
             exclude=libunwind
             enabled=1

             EOF
```

Configuring An Unconnected Network To A Yum Repository

If your system is not connected to a network, you must point the baseurl in the “public-yum-ol7.repo” file to an Oracle Linux 7 ISO DVD image partition.

To configure an unconnected network to a yum repository:

1. Point the baseurl in the “public-yum-ol7.repo” file to an Oracle Linux 7 ISO DVD image partition.

   Note: The following is an example and uses example values only.

```
vi /etc/yum.repos.d/ public-yum-ol7.repo
```

```
[ol7_latest]
  name=Oracle Linux $releasever Latest ($basearch)
  baseurl=http://ap-yummy.us.oracle.com/oracle/OracleLinux/OL7/latest/
    $basearch/
  gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-oracle
  gpgcheck=1
  enabled=0


[ol7_u2_base]
  name=Oracle Linux $releasever Update 2 installation media copy
    ($basearch)
  baseurl=file:///mnt/cdrom/Packages
  gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-oracle
  gpgcheck=1
  enabled=1


  [ol7_UEKR3]
  name=Latest Unbreakable Enterprise Kernel Release 3 for Oracle Linux
    $releasever ($basearch)
  baseurl=http://ap-yummy.us.oracle.com/oracle/OracleLinux/OL7/UEKR3/
    $basearch/
  gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-oracle
  gpgcheck=1
  enabled=0
```

Installing the OS-E Appliance

Once the yum repository is configured, you can install the OS-E appliance onto an Oracle Linux 7 system.

To install the OS-E appliance:

1. Enter the **yum install** command with root permissions.

```
[root@localhost]# yum install asc-appliance
```

2. Reboot your Oracle Linux system so that it starts with the new OS-E appliance kernel required for the OS-E to operate properly.
3. Once you've successfully completed the installation, unmount the OS-E installation USB stick.

```
[root@localhost]# umount   /dev/sdd1 /mnt/usb/
```

Upgrading To 3.7.0M4 From Release 3.7.0M3

This section explains how to upgrade to release 3.7.0M4 from release 3.7.0M3 of the OS-E.

> **Note:** To upgrade to 3.7.0M4 from a release prior to 3.7.0M3, you must first upgrade the OS-E to 3.7.0M3. For information on upgrading to 3.7.0M3, see the *Oracle Communications Application Session Controller Release Notes Release 3.7.0*.

To upgrade the OS-E to 3.7.0M4 you must do the following:

- Backup your OS-E configuration, files, and databases
- Install Oracle Linux 7
- Install 3.7.0M4 from a USB stick
- Restore your configuration, files, and databases on the OS-E

Backing Up the OS-E Configuration, Files, and Databases

Prior to upgrading the OS-E, you must back up your configuration, as well as any necessary files and databases, so that you can restore them once the 3.7.0M4 software is installed.

This upgrade repartitions and reformats the disk, so any items you need (such as .wav files and CDRs) need to be backed up.

To backup the OS-E configuration:

1. Insert your USB stick into the OS-E system to be upgraded.
2. Mount the USB stick using the mount usb command.

```
NNOS-E>mount usb
Device is mounted
```

3. Backup your configuration onto the USB stick using the **restore-stick-create config-backup** command.

```
NNOS-E>restore-stick-create config-backup
```

A folder named “setup”, containing the backed up configuration, is written to the USB stick.

4. Unmount the USB stick using the **unmount usb** command.

```
NNOS-E>unmount usb
```

To backup files and databases:

1. Execute the **database-backup backup system** command.

```
NNOS-E>database-backup backup system <databasePath>
```

2. Repeat this process for as many databases as you need to back up.

For more information on the **database-backup** command, see the *Oracle Communications Application Session Controller Objects and Properties Reference* guide.

Installing Oracle Linux 7.2

Once you have backed up your configuration, you can install Oracle Linux 7.2.

Note: You must install Oracle Linux version 7.2 or higher.

For information on installing Oracle Linux, see Installing Oracle Linux 7.

Installing The OS-E

When Oracle Linux 7 is installed, you can download the OS-E installation file, copy it to a USB stick, and install the OS-E.

For more information on installing the OS-E, see Installing the OS-E.

Restoring Your Configuration, Files, and Databases On The OS-E

When the OS-E is installed, restore your configuration, files, and databases.

To restore your OS-E configuration:

1. When the OS-E has rebooted and the login prompt appears, insert the USB stick you created in Backing Up the OS-E Configuration, Files, and Databases.
2. Mount the USB using the following command:

```
NNOS-E>mount usb
```

3. Execute the **restart warm** command.

The configuration is loaded onto the OS-E.

To restore your files and databases:

1. Execute the **file fetch** command or SCP to copy any .wav files and archived CDRs to the OS-E.
2. Execute the **database-backup restore system** command to restore any previously backed up CDRs and databases.
3. Repeat this process for as many databases as you need to restore.

For more information on the **file fetch** and **database-backup** commands, see the *Oracle Communications Application Session Controller Objects and Properties Reference* guide.

Chapter 3. Quick Commissioning New OS-E Systems

About This Chapter

The chapter provides the basic information that allows you to configure OS-E software after you have physically installed the system in your network. Commissioning enables an OS-E system or compatible third-party device to process locally registered SIP phone calls.

Prerequisites to Quick Commissioning

Before using the information in this chapter, make sure that you have properly installed and cabled the system, as covered in Chapter 1. The following OS-E documents provide additional information on configuring Session OS-E services, as well as how manage the system using the OS-E CLI and the OS-E Management System.

- *Net-Net OS-E – System Administration Guide*
- *Net-Net OS-E – Objects and Properties Reference*
- *Net-Net OS-E – Management Tools*

Additionally, the *Net-Net OS-E – Release Notes* provides important information about the software that you should review before commissioning a system in your network.

Steps 1 through 5 cover the tasks and services for getting the system up and running on an IP network so that the Ethernet interfaces can process SIP sessions. When enabled on an IP network, you can manage the system and its configuration remotely over the Internet using the OS-E Management System.

Steps 6 through 10 cover the tasks that allow you to control and monitor SIP sessions, as well as store call detail records and recordings.

Building the Configuration File

The OS-E configuration file (*cxc.cfg*) is made up of configuration objects and property settings that control how the system processes and manages SIP traffic. As you open these objects and set properties using the CLI or the OS-E Management System, the software builds a configuration hierarchy of objects that are applied to SIP sessions. You can display this configuration hierarchy using the **show** and **show -v** (verbose) commands.

For new users, as well as for users who are adding functionality to their configuration, you will need to open configuration objects using the **config** command to enable the default settings for those objects, even if you choose not to edit any of their associated properties. For example, if you need to enable the **ICMP** protocol and its default settings, you simply open the object and execute **return**, as shown in the session below. Notice that the ICMP object has been added to the configuration hierarchy at the end of the session on the eth4 interface.

```
config> config box interface eth4
config interface eth4> config ip 172.26.2.14
config ip 172.26.2.14> config icmp
config ip 172.26.2.14> return
config interface eth4> return
config box> return
config> show -v
interface eth4
  admin enabled
  mtu 1500
  arp enabled
  speed 1Gb
  duplex full
  autoneg enabled
  ip 172.26.2.14
   admin enabled
   ip-address dhcp
   geolocation 0
   metric 1
   classification-tag
   security-domain
   address-scope
   filter-intf disabled
   icmp
    admin enabled
```

```
limit 10 5
```

To remove an object from the configuration hierarchy, use the CLI or OS-E Management System **delete** command.

Basic Network Topology

The following image illustrates a network topology using the OS-E with a directly-attached PC for initial setup, and the OS-E Management System for remote access using a graphical user interface.

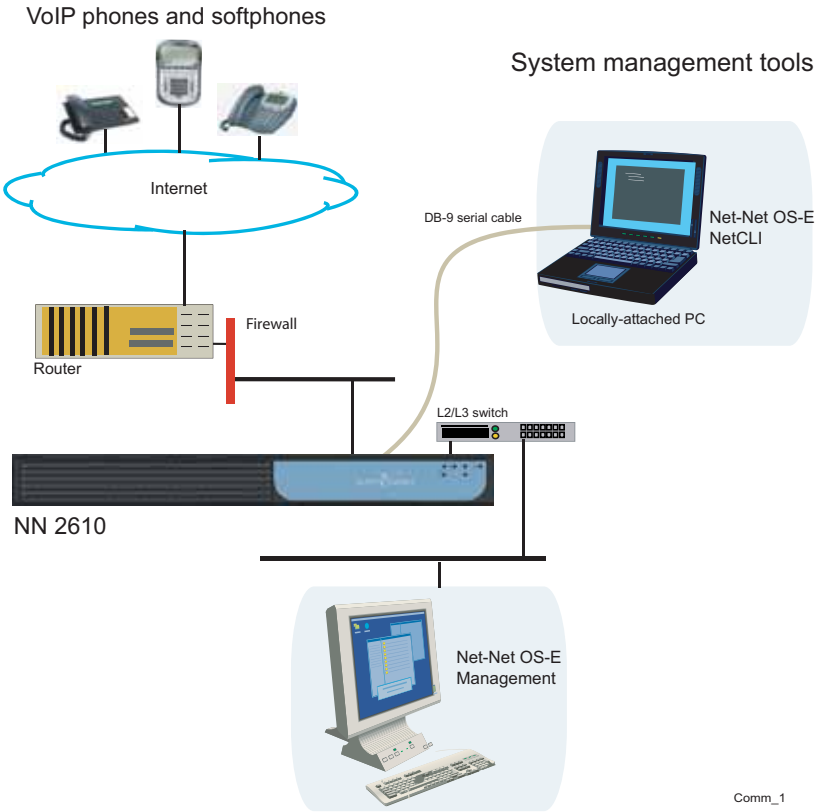

Step 1. Configuring Basic IP Connectivity

Before you can manage an OS-E system remotely over the Internet using the OS-E Management System or over a Telnet or SSH connection, you need to locally assign an IP address to one of the Ethernet interfaces, **eth0**, **eth1, eth2,** or **eth3.** If you are setting up the device remotely, you will also need to configure an IP route, a route to a destination host or network, and a gateway IP address.

If you are using the OS-E Management System, you will also need to know the assigned IP address on one of the Ethernet ports to manage the OS-E configuration. The OS-E Management System application runs directly on the OS-E system over the Internet using the Internet Explorer Web browser.

The following CLI session creates and enables an IP interface named **192.168.124.5,** sets the static IP address and network mask, configures an IP route (if connecting remotely), and enables Web access on this IP interface. You will need to enable ICMP on the OS-E IP interface before you can use the **ping** command from your console to test the device as a responding node on the network. Use the **show -v** command to display the configuration.

CLI Session

```
NNOS-E> config box
config box> set hostname local2610
config box> config interface eth1
config interface eth1> config ip mgmt-int
Creating ‘mgmt-int’
config mgmt-int> set admin enabled
config mgmt-int> set ip-address static 192.168.124.5/24
config mgmt-int> config routing
config routing> config route internetGateway
Creating ‘route internetGateway’
config route internetGateway> set destination default
config route internetGateway> set gateway 192.168.124.3
config route internetGateway> return
config routing> return
config ip mgmt-int> config web
config web> set admin enabled
config web> set port 80
config web> return
config mgmt-int> config icmp
config icmp> set admin enabled
config icmp> top
config> save
config> show -v
```

Using the Setup Script

An optional configuration setup script called *cxc.setup* is now included with newly shipped systems. After installing a new system, you can run the script directly from the NNOS-E> prompt, as shown in the example session below.

The script presents a set of questions to help you with the initial system configuration. The information in the script includes the following:

- Local hostname
- IP interface names and addresses
- SSH and Web access
- Default route and any additional static routes per interface for remote management
- User-defined CLI prompt

Every OS-E system has a minimum of two Ethernet interfaces. Any Ethernet interface on the system can be used for management traffic, however, Oracle recommends the use of eth1, as eth0 is reserved for fault-tolerant clustering with other OS-E systems. Management traffic is also supported on any interface that is carrying private or public network traffic. This means that it would be possible to use eth1 to carry SIP traffic and management traffic.

CLI Session

```
NNOS-E> config setup
set box\hostname: <name>
config box\interface: eth1
set box\interface eth1\ip a\ip-address: <ipAddress/mask>
config box\interface eth1\ip a\ssh (y or n)? n
config box\interface eth1\ip a\web (y or n)? y
config box\interface eth1\ip a\routing\route: <routeName>
set box\interface eth1\ip a\routing\route localGateway\gateway:
<ipAddress>
set box\cli\prompt: <newPrompt>
Do you want to commit this setup script (y or n) y
Do you want to update the startup configuration (y or n)? y
```

Note: The /cxc directory on the OS-E system may include vendor-specific scripts that address unique startup configuration requirements. Specify the name of the script on the command line following the **config setup** command. For example:

```
NNOS-E> config setup vendor.setup
```

Check the /cxc directory for any vendor-specific setup files included with your system.

Enabling Network Access

To ensure you can manage the system using services such as Telnet or the OS-E Management System, you must configure the OS-E system so that it is available on the network. You need to create a default (or static) IP route, a route to a destination host or network, and a gateway IP address.

After you configure the static route, enable ICMP and then use the **ping** command at the top-level of the CLI to test network accessibility.

Defining a Default Route and Gateway IP

If you are setting the box remotely, you will need to configure an IP route, a route to a destination host or network, and a gateway IP address.

Refer to Step 1. Configuring Basic IP Connectivitythe previous section in this chapter, for the example CLI session that shows the routing context and the route named *internetGateway*. This is the default route that uses 192.168.124.3 as the default gateway.

Launching the OS-E Management System

In addition to the CLI, you can use the OS-E Management System to configure the OS-E. To access the OS-E using the OS-E Management System, open an HTTP or secure HTTP window (HTTPS) to the IP address of the Eth0 port on the OS-E system. For example:

https://192.168.124.5

You should see the Oracle OS-E Log In window, illustrated in the following image.

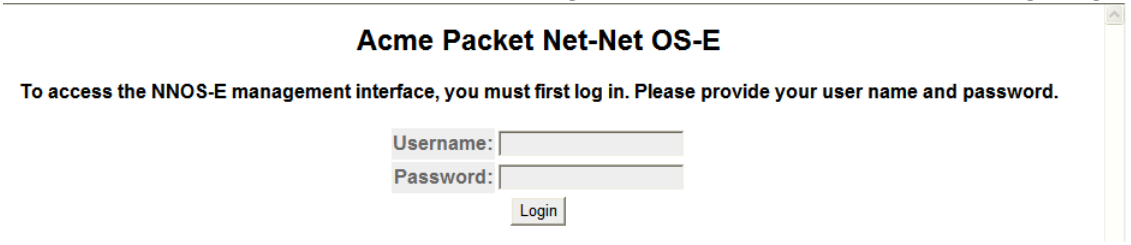

By default, there are no user accounts configured on a new system. This means any value can be entered in for username & password, or leave the fields blank and click **Login**. Once you log in, the OS-E Management System main page appears.

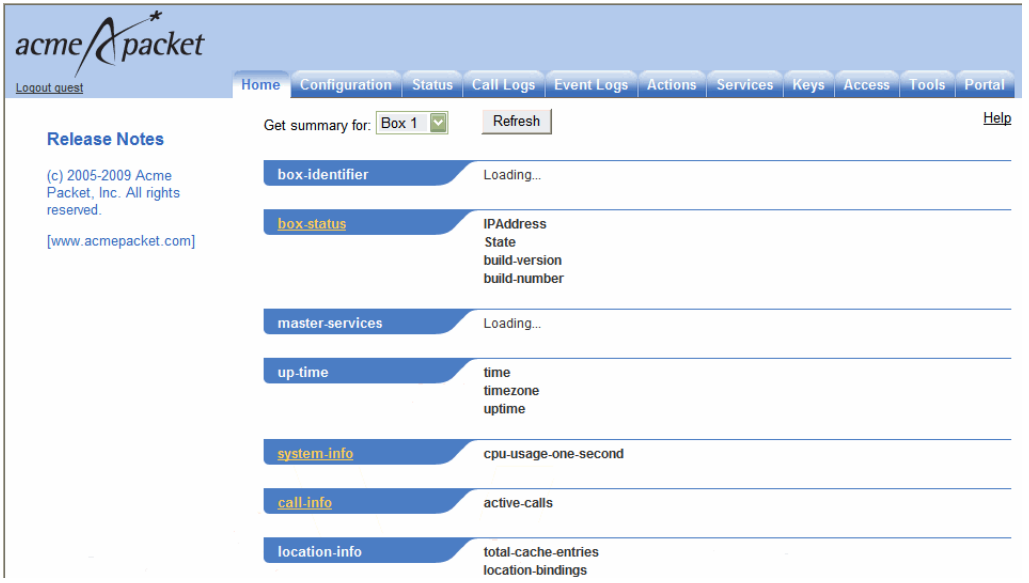

The remaining steps in this chapter use the OS-E Management System to commission the OS-E.

Changing the Linux Root Password

To change the Linux root password, use the secret root action. When prompted, specify and confirm the new password. For example:

```
NNOS-E>secret root
password:*******
confirm:*******
```

```
Success!
NNOS-E>
```

Note: The password must be at least four characters long.

For more information on the **secret root** action, see the *Oracle Communications Application Session Controller Objects and Properties Reference Guide*.

Step 2. Configuring Advanced IP Connectivity

Use the **Configuration** tab or the CLI to configure several additional Ethernet interfaces, as covered in Step 1. As a security device, the NN 2600 Series uses a default setting of **disabled** for these objects in the configuration file. This means that you must enable each interface. These objects include:

- **SSH**—To enable SSH client connectivity on the interface
- **Media ports**—To enable a range of port numbers for on the interface
- **SIP**—To enable SIP traffic on the interface)

When editing Ethernet interface and examining each object using the OS-E Management System, note that many of the objects are already visible, but they are not yet enabled. For these objects to actually be enabled on the OS-E system, you must select the object and save the configuration.

After editing an interface configuration, elect **Set**, then **Update & save configuration**, as illustrated in the following image.

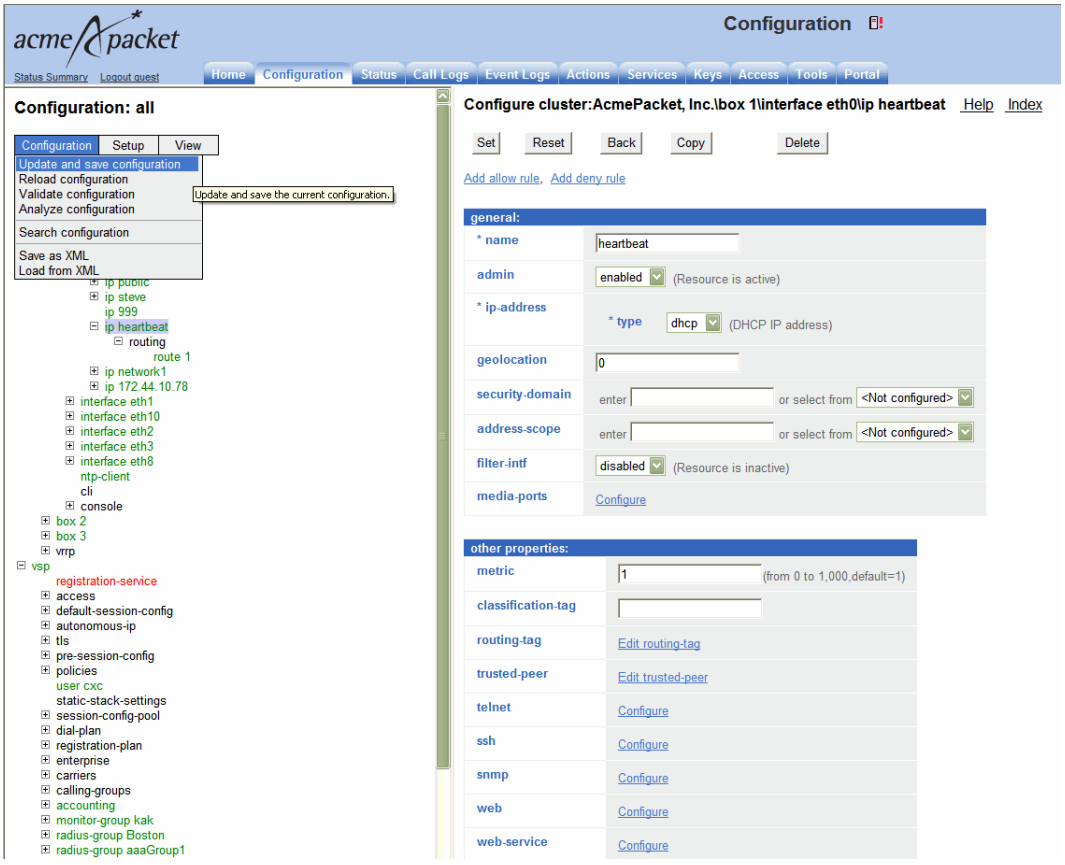

When you select **Configuration/Update and save configuration** you will be asked "Do you want to update the live configuration?" followed by "Do you also want to save the live configuration?" Click **OK** for both questions to ensure that the configuration is properly saved to the OS-E configuration file, *cxc.cfg*.

The following steps are necessary to set some specific parameters for the objects listed above:

1. Select the Configuration **Cluster/ Box 1/Interface Eth0/IP local** object on the left menu tree. Under the **General** field, edit the Media Ports properties as desired, then click **Set**.
2. Under the **Other Properties** field, edit the SSH properties. Accept the defaults by clicking **Set**.

3. Select **SIP** from the menu tree. Enter the following values for each fields:
   - admin: enabled (default)
   - NAT translation: disabled (default)
   - UDP port: Select **Add UDP port**, accept the defaults, then click **Finish**->**Set**.
   - TCP port: Select **Add UDP port**, accept the defaults, then click **Finish**->**Set**.
   - TLS port: Select **Add UDP port**, accept the defaults, then click **Finish**->**Set**.
   - Certificate: blank (default)

When you are finished editing the SIP fields, select **Set->Configuration/Update and save configuration**.

Step 3. Creating User Accounts for Basic Access

By default, the OS-E does not contain any predefined user accounts. This means it is possible to access the management interfaces without entering any login credentials (username and password). You are not required to create user accounts, but it may be desirable for security reasons. If you want to create a user account at this time, follow the steps below. If not, go directly to Step 4.

1. Using the OS-E Management System, select the **Access** tab, then select **Access** from the left menu pane. The Access Permissions/Configure Access page appears.

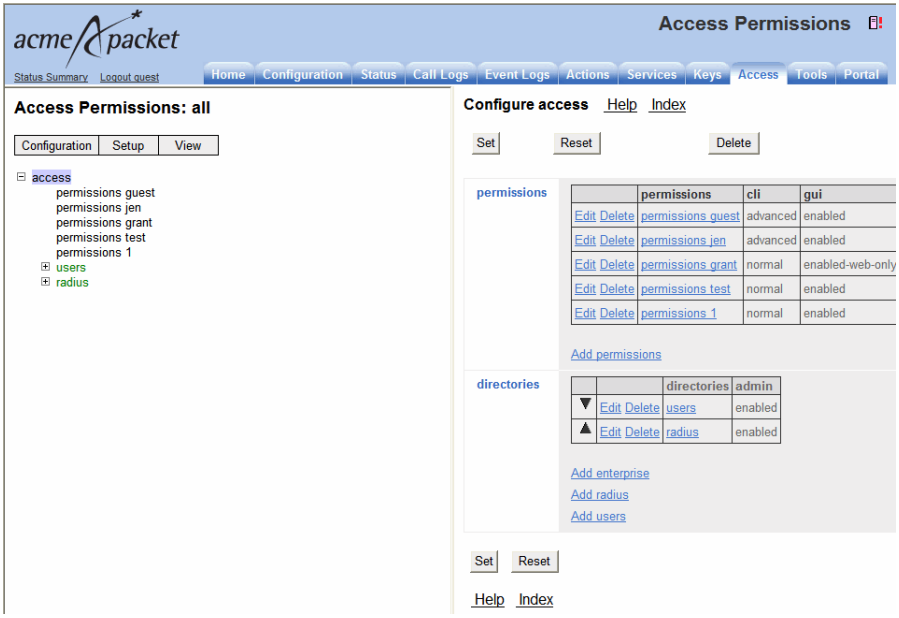

2. Under **permissions**, select **Add permissions** and create a permissions group called *super-user* and accept all default settings with all permission types enabled.

   Select **Set**, then select **Update and save configuration** from the Configuration pull-down in the left pane**.**

3. From the **Directories** object, select **Add users**. Accept the default setting of enabled.

4. Select **Add user** and enter the required **name** and **password** of your choice, then re-enter the password to **confirm** your original password entry. In the **permissions** field, choose the permissions group that you just created (super-user).

5. Click **Create.** Select **Configuration->Update and save configuration**.

These steps created a username and password for a super-user account. Future attempts to log in to the OS-E (using the CLI or the OS-E Management System) will require that you specify these login credentials. If needed, you can also create user accounts with one or more of the super-user permissions.

Step 4. Enabling Master Services

The **master-services** configuration enables directory, accounting, database and registration services to run on the system. Perform the following steps to configure these master services:

1. Select the **Services** tab, then select **master-services** form the left menu pane.
2. Accept the default settings for **cluster-master**, **directory, accounting**, **database** (with **Show advanced** button selected), and **registration**. Click **Set**.

After you have configured all five services, select Configuration->**Update and save configuration**. The completed Master Services configuration should appear as shown in the following image.

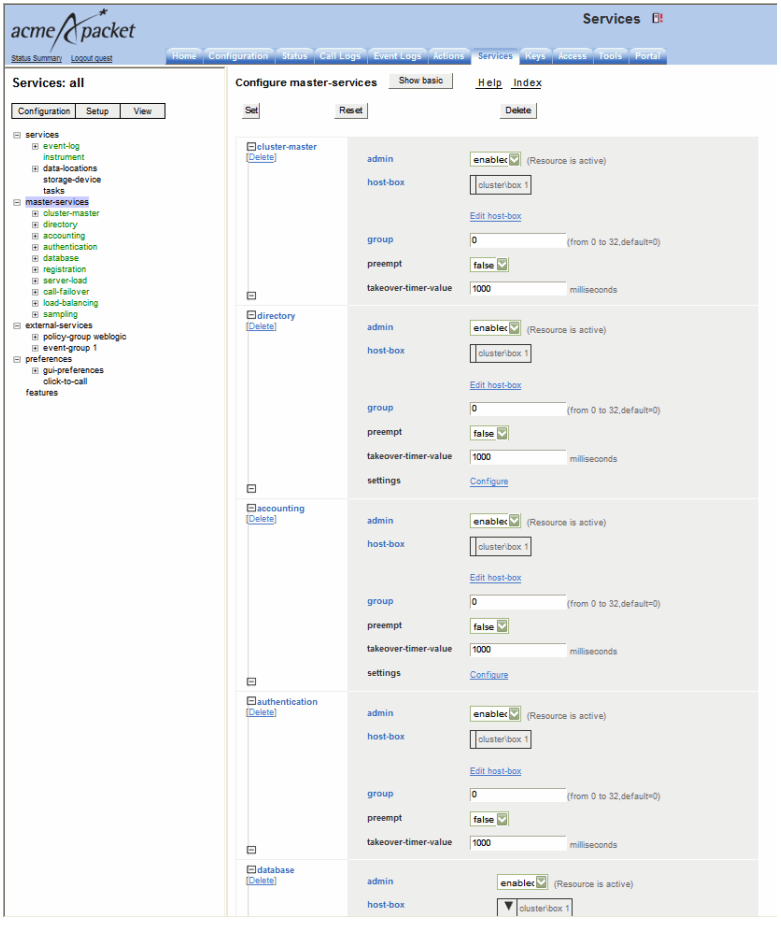

Step 5. Configuring Basic Services

The **Services** configuration enables event logging and virus scanning services to run on the OS-E. Perform the following steps to configure event logging on the system.

1. Select the **Services** tab then select Services from the left menu pane.
2. On the Configure services page, select **event-log** from the menu pane, accept the defaults and click **Set**. Under the **event-log** configuration, additional options are available that you can configure, as illustrated in the following image.

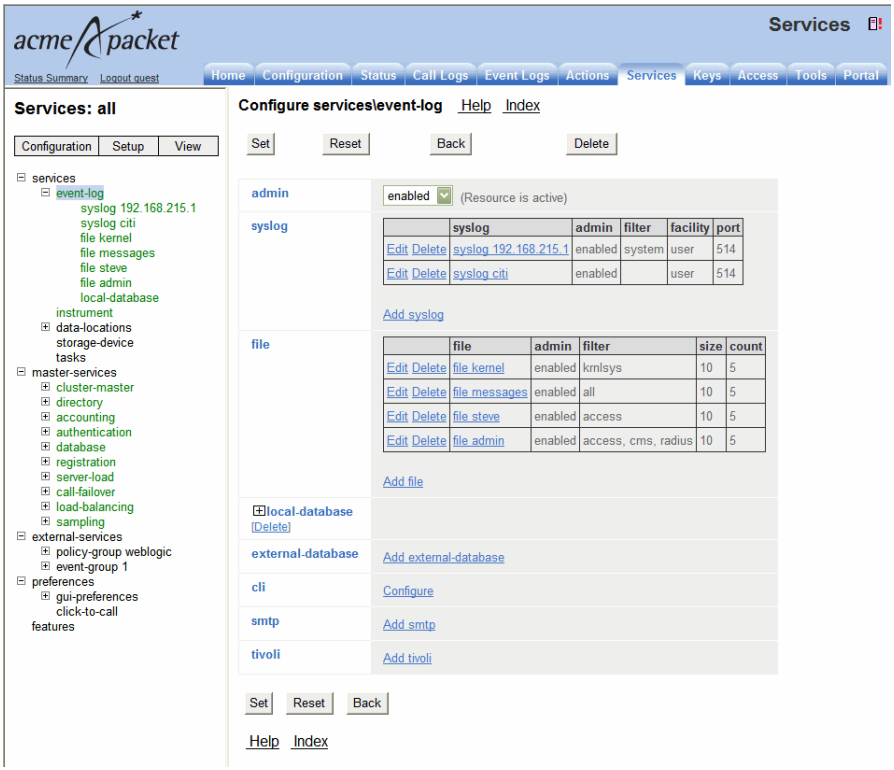

You can direct the event logs to one or more of the following locations:

- A syslog server
- An ASCII file in an OS-E directory
- A database on the OS-E system
- An external database

The following image shows a configuration that specifies that logs should be directed to a syslog host (at 192.168.215.1), a local file on the system, and the local database. The syslog system will receive messages of the system severity (or lower). The local file is named *messages* is created in the log directory.

The configuration also shows two filters: the first filter captures events of the *system* class with **debug** severity level, and the second filter captures event messages that match the **error** severity level. Refer to the *Net-Net OS-E – System Administration Guide* for information about event logs, syslog, and event filters.

3. In the **file** object, click **Edit**, then enter the name *event-log* in the text block. Click **Set**.

   This configures event logging so that messages are written to the local file named *event-log*.

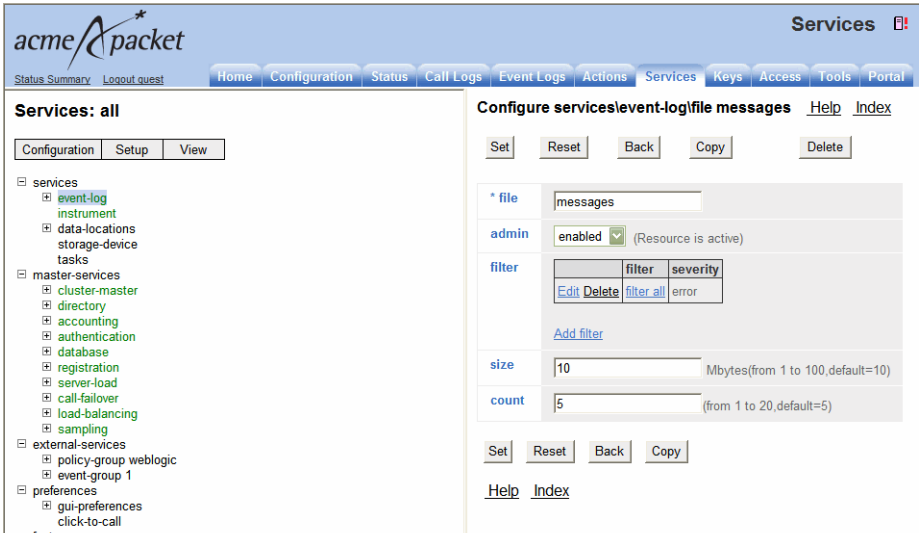

Step 6. Enabling the Virtual System Partition (VSP)

The OS-E virtual system partition (VSP) is the part of the system that holds the comprehensive customer-defined configuration that controls how the system processes, stores, directs, and routes SIP traffic. The VSP is where you can create session configurations, registration and dial plans, and policies that handle SIP REGISTER and SIP INVITE traffic (and other SIP methods) that the system will receive and forward to a SIP call destination, authentication and accounting database, VoIP service provider or carrier, enterprise server, and so on.

Using the OS-E Management System, perform the following steps.

1. Select the **Configuration** tab, then select **vsp** from the menu to open the Configure vsp page, as illustrated in the following image.

2. Under the general heading:, change the **admin** state to **enabled**.
3. Click **Set,** then select **Configuration->Update and save configuration**.

Step 7. Configuring the Accounting Environments

This step is necessary to configure the system to store call detail records and voice call recordings.

1. Select the OS-E Management System **Configuration** tab, then select **vsp->accounting** from the menu to display the Configure vsp\accounting page, as illustrated in the following image.

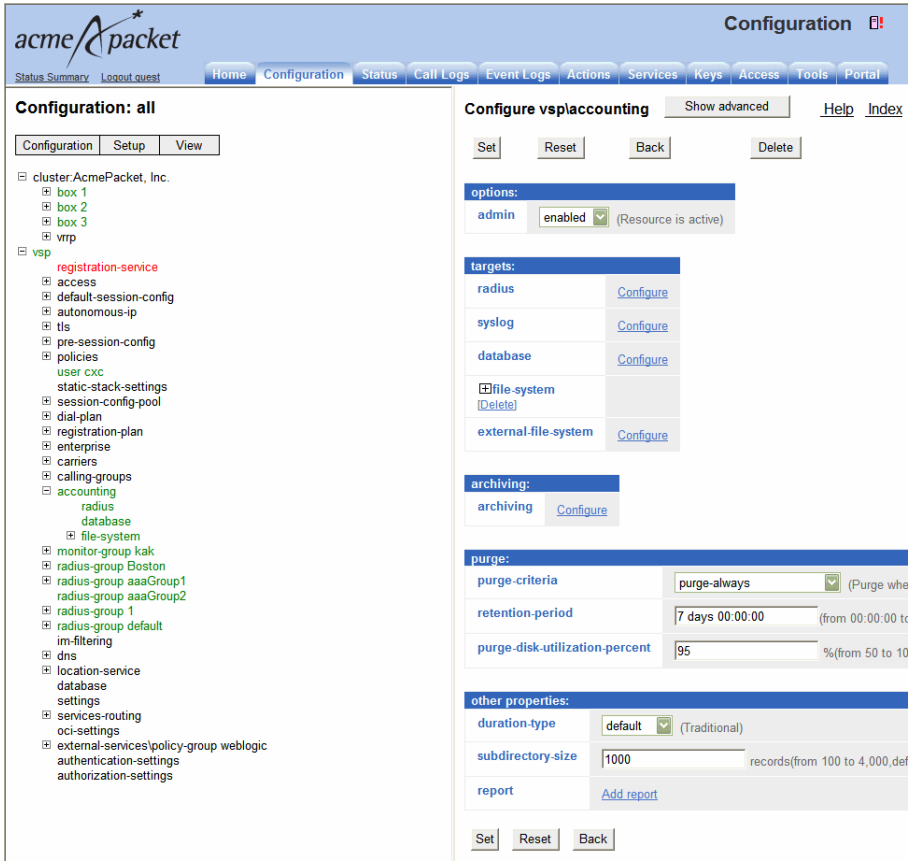

2. Under **targets**, go to the **database** and set the **admin** property to enabled.
3. Select the database **Add group** command. The Edit group screen appears, as illustrated in the following image.

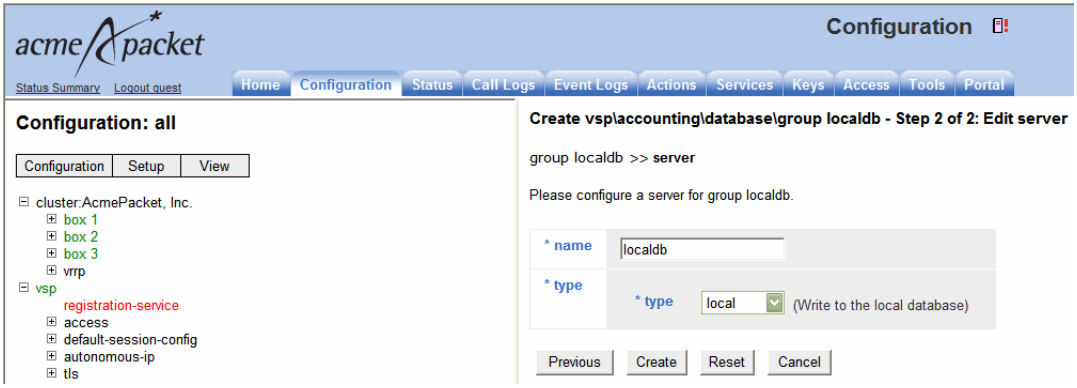

4. Enter *localdb* in the **target-name** field and select **Create** to display the Configure database group page**,** as illustrated in the following image..

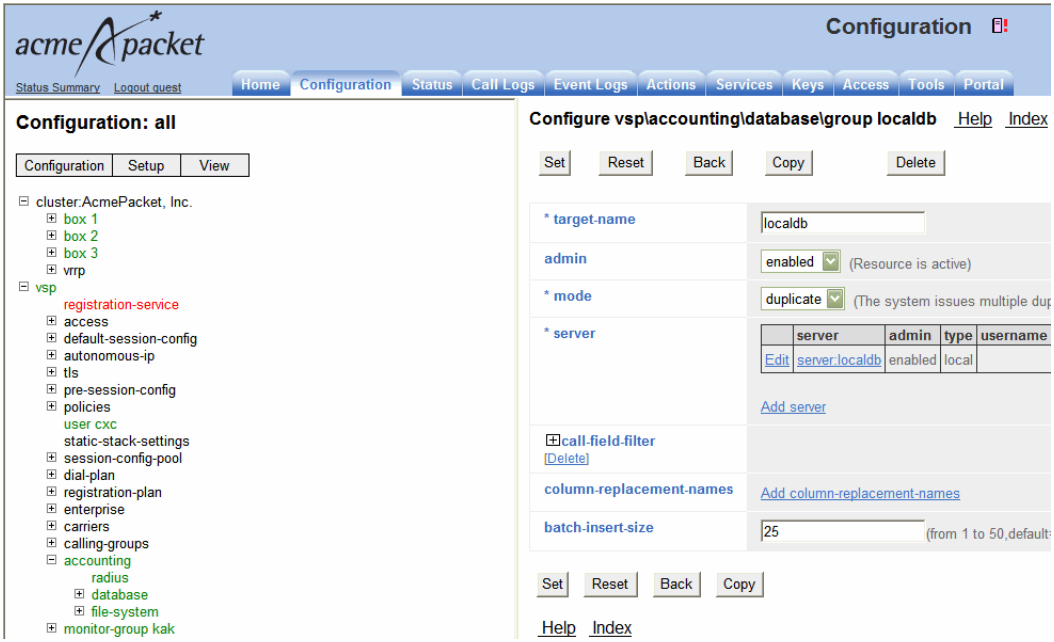

5. Click **Edit** and configure the following settings:
   - **admin:** enabled
   - **name:** localdb
   - **type:** Select **local**

- **username**: postgres
- **password-tag:** postgres

Note: If you set the server **type** to *local*, using the local database as the accounting target, set the **username** and the **password-tag** to *postgres*. If you edit the **username** and **password-tag** properties to anything other than *postgres*, data will not be written to the database.

For information about password tags, refer to the *Net-Net OS-E – Objects and Properties Reference*.

6. Click **Set,** then select **Configuration->Update and save configuration**. The screen appears as illustrated in the following image.

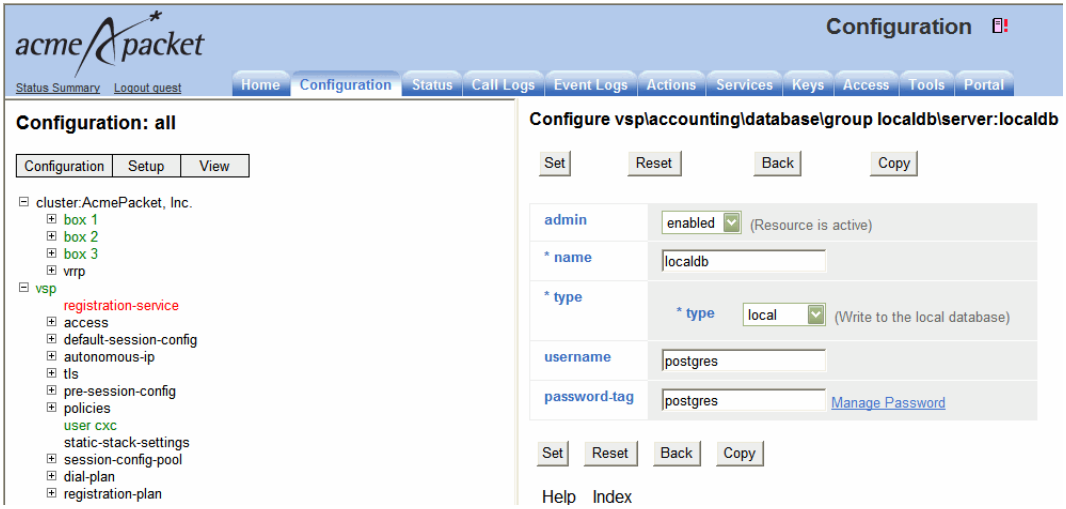

Step 8. Editing the Default Session Configuration

Step 8 configures a default system policy that allows the OS-E to process SIP traffic. By default, and for security purposes, the OS-E does not allow any SIP traffic to pass.

1. Select the Configuration tab, then select **vsp->default-session-config** from the menu to display the vsp/default-session-config page, as illustrated in the following image (top portion).

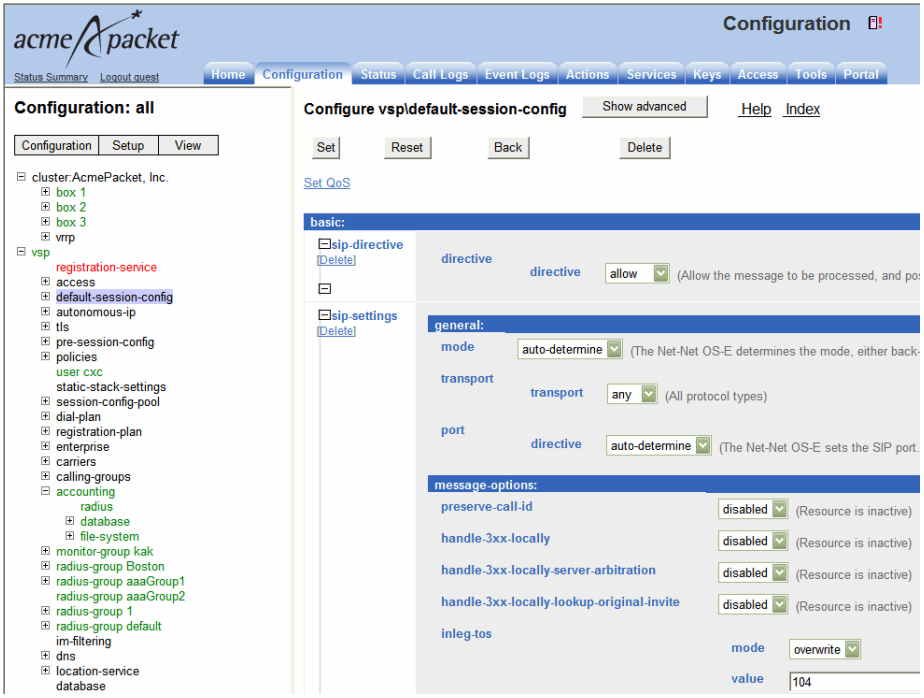

2. In the **sip-directive** object, change the directive policy to **allow**, if not already set. This allows SIP traffic to traverse the OS-E system. Click **Set**.
3. Scroll down to the **media** object. Change the **anchor** and the **recording-policy/ record** properties to **enabled.** Accept all other default settings.
4. Click **Set**, then select **Configuration->Update and save configuration.**

Step 9. Enabling Registration Services

Step 9 enables the OS-E to handle SIP REGISTER sessions, allowing locally registered SIP clients to pass SIP sessions, as well as forward REGISTER sessions to upstream destination registrars.

1. Select the Configuration tab, then select **vsp/ registration-service** from the menu tree to display the vsp/registration-service page, as illustrated in the following image.

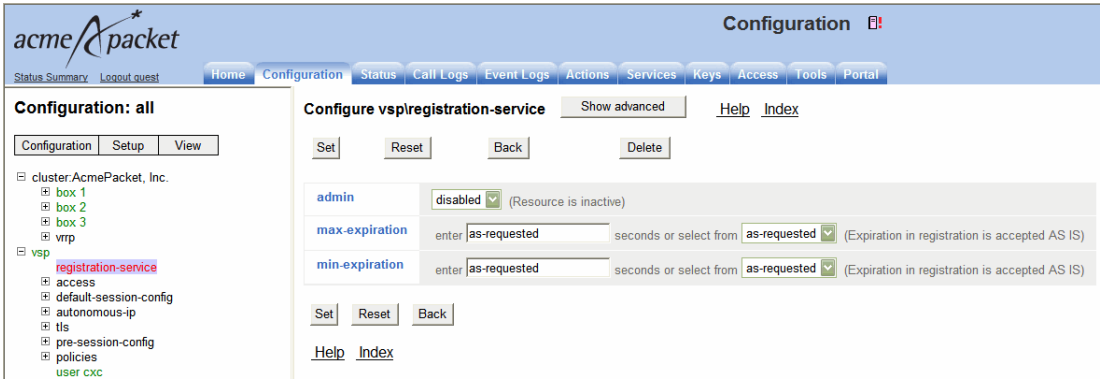

2. Accept the default settings and click **Set**. This enables the registration service on the OS-E system.

The OS-E will now provide support for basic SIP calls between locally registered clients.

Step 10. Reviewing the Configuration

Once you have completed Steps 1 though 9, review the configuration to make sure it is accurate. A quick way to do this is to scan the OS-E Management System navigation tree to make sure there is an entry for each of the objects that you configured.

The following image is a listing of the Configuration and Services objects configured as part of basic OS-E commissioning. If you are using the CLI, run the **show -v** command from the NNOS-E prompt to display the configuration that you just created. The following image displays the configuration and services navigation trees.

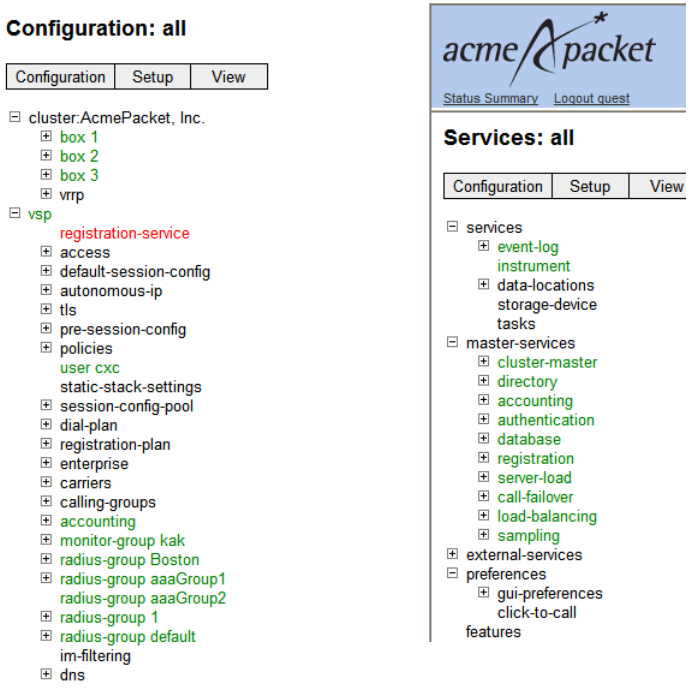

Chapter 4. Installing OS-E Clusters

About This Chapter

This chapter provides information on how to install an OS-E cluster, a group of OS-E systems that operate together to support redundancy and failover, high-availability, load balancing, and configuration.

OS-E Cluster Overview

A “high-availability” cluster is a group of OS-E systems that provides a single point of configuration management, and at the same time, expands functionality across multiple devices participating in the cluster. An OS-E *master* manages the configuration for the entire cluster. All members of the cluster share network resources, network load, media ports and streaming, registration, and other processes.

OS-E systems within a cluster may be geographically dispersed in the network. A cluster recovers from the failure of one or more cluster members through health monitoring, shared master services migration, and network redundancy using the Virtual Router Redundancy Protocol (VRRP).

A cluster can be set up to operate as a two-system primary/standby redundant configuration.

Cluster Operations and Services

In the two-system redundant configuration, one OS-E system is the active master, performing signaling & media processing, and the other OS-E system is available as a standby system for the signaling & media processing if the master fails. Master failover allows another OS-E system to assume the master role in the cluster should the originally configured master become unavailable. VRRP is responsible for handling the failover from the master to the backup device.

Master-Services

The **master-services** configuration is responsible for mirroring the state of the cluster to allow reliable failover to a standby device. The following sections describe the suggested settings for the **master-services** objects:

Cluster-Master

A **cluster-master** configuration on the OS-E system designated as the master is responsible for passing configuration changes to cluster members. A secondary property called **takeover-timer-value** specifies the number of milliseconds (such as 500) that the master-service stays in “awaiting takeover” mode at boot time.

Use the **show -v** command to display the **current takeover-timer** value. When the OS-E boots, each hosted master-service waits for this period to determine if any existing devices in the cluster are already running that service before assuming mastership.

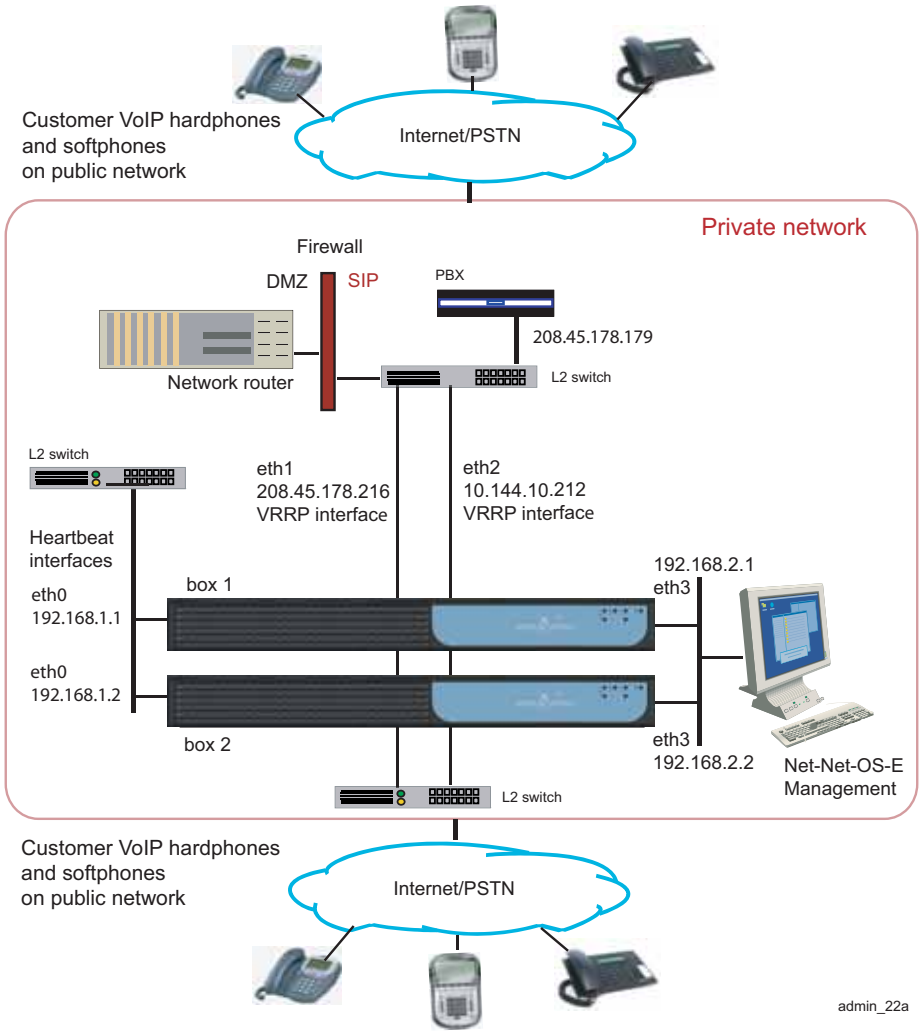

Directory

When enabled, directory services allows the OS-E master to use enterprise (or corporate) directories that contain the identities of SIP users who are authorized to access the SIP enterprise communications servers.

In environments running CSTA, it is necessary to link the **directory** master service into a VRRP group with an interface to reach a Broadsoft OCI server. If the directory service is running on an OS-E system that cannot reach the OCI server, the CSTA-to-OCI translation will not function. Use a **takeover-timer-value** of 500 milliseconds.

Accounting

When enabled, accounting services supports RADIUS accounting, system logging (syslog), DIAMETER protocol services, the accounting database, and the accounting file-system

Authentication

Authentication services enables or disables all authentication functions on the OS-E, such as RADIUS and local user profiles. If authentication is disabled, you can still configure the authentication services, but the services do not become active until you enable this master service.

It may be necessary to link the authentication service to VRRP interface(s) using a group configuration if the VRRP interface is used to contact the authentication servers. Use a **takeover-timer-value** of 500 milliseconds for authentication.

Database

The master-services **database** object allows you to configure maintenance and other settings for the OS-E system database. The OS-E database is the local repository for call accounting records and media files

The **database** master service should be on a backup OS-E system, with the secondary property **preempt** set to *true*. This will help maintain the data in one location in the event of a brief service outage.

The **preempt** property specifies whether the master-service should resume the mastership if it has gone down and then returned to operation. If set to *true*, the master resumes its position. If set to *false*, the backup service retains master control.

Server-Load

The master-services **server-load** object configures the OS-E to calculate server load. This object must be enabled if your dial plan arbiter rule settings use **least-load** as the routing algorithm option. (The arbiter rule property sets the criteria by which the OS-E selects the server to which it forwards calls.)

Configure the **server-load** master-service for outbound server load balancing or server based admission\emission control. Currently, the **server-load** master-service should be linked to the VRRP SIP signaling interfaces over a configured group.

Call-Failover

The **call-failover** master-service configures failover for the media and signaling streams. As a master-service, the configured host OS-E master distributes copies of the media and kernel rules to all backup devices in a cluster. The OS-E uses the database on the host box, but enabling **call-failover** ensures that there is an active copy of the database on another device in the cluster in the event of a failure.

Registration

Enabling the registration service allows the OS-E to accept SIP REGISTER requests in behalf of other SIP servers (called *registrar peers*) that reside in other domains.

The **registration** master-service configures the registration process for intracluster registration lookups. In a cluster, the registration database runs on the specified master and the selected backups. The **host-box** property establishes the master and selective mirroring. The first OS-E listed is the master, while subsequent devices have mirrored databases. The OS-E systems not configured with the **host-box** property use the local location cache instead of the registration database. The **registration** master-service must be enabled for load-balancing of SIP processing (across backing interfaces configured with the **sip** object) to function correctly.

Load-Balancing

The master-services **load-balancing** object configures OS-E systems to host the load-balancing master service. For detailed information, see Configuring Cluster Load Balancing.

File-Mirror

The master-services **file-mirror** object sets all participating OS-E systems to share particular files (the types of files shared are preset in the OS-E operating system), such as media recordings, log files, etc. The file-mirror master service distributes files to all OS-E systems listed as hosts for the service.

Once the files are mirrored, you can play them back from any OS-E system that functions as a host.

Least-Cost-Routing

The master-services **route-server** object sets the route-server master service, which manages the server process. The master service handles requests from local or remote OS-E systems for least cost route definitions.

For detailed information on the route-server, see the following manuals:

- *Net-Net OS-E – System Administration Guide*
- *Net-Net OS-E – Objects and Properties Reference*
- *Net-Net OS-E – Session Services Configuration Guide*

Sampling

The master-services **sampling** object opens the mechanism for setting the interval at which the OS-E samples operational aspects of the system for either:

- Display in the OS-E Management System, or
- For sending to an IBM Tivoli server

By setting sampling for a status provider, you can view data for that provider over a specified period of time. The OS-E supports two sampling targets—a Postgres SQL database and an IBM tivoli server. (Set the provider data sent to the target using the **status** and **provider** objects. See *Net-Net OS-E – Objects and Properties Reference* for more information on configuring these objects.)

Once you have enabled **sampling**, the master service stores the samples in its local database.

Third-Party-Call-Control (3PCC)

The master-services **3pcc** (third-party-call-control) object configures call control, allowing the OS-E or a CSTA client to control (become the third party) in a call. Specifically, this object controls the WAV files that the OS-E should play and the external status events reported to an external server for calls created by the OS-E.

For detailed information on CSTA, see the following manuals:

- *Net-Net OS-E – System Administration Guide*
- *Net-Net OS-E – Objects and Properties Reference*
- *Net-Net OS-E – Session Services Configuration Guide*

Heartbeat Interface, BOOTP, and Messaging

Use the Ethernet physical interface **eth0** as the heartbeat interface for the OS-E cluster. This interface is used by default for any backup OS-E system that you added to the cluster. The systems will perform a BOOTP request over that interface and you will be able to add these systems by creating an entry for each to the configuration, and then booting them.

Once an OS-E is a member of the cluster, that system will receive a saved configuration file (*cxc.cfg)* from the master. Each time the *cxc.cfg* file is saved on the master, the latest copy of the cxc.cfg file is sent to each device in the cluster. You will need to configure a messaging interface on each cluster member so that the master knows the interface over which members of the cluster will receive the cxc.cfg file.

Event Logging

Event logs are stored on each box individually and represent the events that occurred on that particular OS-E system. You configure event logging in the **services/event-logs** configuration object. The recommended event log filters on a cluster are as follows:

- Local-database all error
- File *name* system error
- File *name* krnlsys info
- File *name* system info

- File *name* db info

Network Time Protocol (NTP)

Ensure that you have NTP configured on all OS-E systems, ensuring that they point to a timeserver which will keep their time synchronized. DO NOT use a VRRP interface as your route to the timeserver, since one device will always have the VRRP interfaces down and will not be able to contact the NTP server.

If you do not have access to an external NTP server, configure one of the clustered OS-E systems to be an NTP server for the other cluster members. It is important to run NTP, as the time on all clustered system must be kept synchronized. If the times on the OS-E systems drift apart, the Denial of Service (DOS) software will not function properly, as timestamps are required to make this work across the cluster.

You can configure the NTP-server on the messaging interface on one OS-E system, and have all other devices point to this IP address in their NTP-client configuration.

Cluster Redundancy Operations

The OS-E cluster redundancy operates as follows:

- Internal messaging is exchanged so that each OS-E system knows the state of the other boxes, either up or down.
- If the active cluster master goes down, the box listed next in the list of cluster masters becomes the active cluster master. (Note that mastership does not automatically go back to the original system when it returns to service.)
- All the other master services work similarly, with an ordered list of devices that can run the service and the active service running on the next device in the list if the active master fails.

If an OS-E system fails, another device in the cluster will assume its network interfaces using VRRP.

Notes on Cluster Management

The OS-E cluster management operates as follows:

- Within a given cluster, one box functions as the active cluster master.

- Configuration and management of all boxes within a cluster is performed through the cluster master.
- There are no limitations on how many boxes within the cluster can be configured as backup cluster masters or backups for any of the master services.
- The configuration contains a list of boxes that can be cluster masters. The ordering of this list reflects the order in which boxes attempt to become master (i.e. the box listed first becomes the initial master, if that box fails then the next box in the list attempts to become the master, etc.)
- The OS-E Management System connects to the cluster master and provides a single point of management for the following:
  - Configuration
  - Status reports
  - Call logs
  - Accounting data
  - Actions
- The CLI provides single point of management for configuration using the CLI on the cluster master. The CLI is still available on all the other devices in the cluster, so any CLI commands can be executed on individual boxes.
- Note that the management functionality available from a given cluster is dependent on the functionality being performed by that cluster. For example, call logs are available only on clusters where signaling is performed; media recordings are available only on clusters where media streaming is performed.

Cluster Installation Prerequisites

Before beginning the cluster installation, ensure that any L2/L3 switch supporting the cluster has the Port Fast, Fast Link, or similar feature turned on. This allows the switch to run the Spanning Tree 802.1 protocol so that the switch ports being used by the OS-E go directly to the “forwarding” state. If the switch does not support Port Fast or Fast Link, disable the Spanning Tree protocol for the VLANs associated with the switch ports being used by the OS-E.

Cluster Installation Procedure

There are a number of steps that you need to follow to install an OS-E network cluster. You will need to know certain information about all the systems in the cluster for proper operation.

Each step uses a sample CLI session of commands that best illustrate how to best configure important settings.

1. Determine the specific OS-E system to assume the role of cluster master. Configure **master-services** to specify the device the cluster to assume initial mastership.

```
NNOS-E> config master-services
config master-services> config cluster-master
config cluster-master> set admin enabled
config cluster-master> set host-box cluster\box 1
config cluster-master> set host-box cluster\box 2
config cluster-master> set group 1
config cluster-master> return
```

2. Note the MAC address (identifier) on each device in the cluster. The MAC address is on a sticker on the back of the system. Write down each MAC address on a pad or piece of paper.

   On each device, if there is no sticker present, attach a laptop or standard PC to the system console port and perform the following steps:

   — Power up the system

   — At the NNOS-E prompt, execute the **show interface-details eth0** command to display the MAC address.

3. Attach a console to the cluster master and power up the OS-E system.

4. Configure the cluster master by configuring the Ethernet interfaces, IP addresses, and protocols. Ethernet interface eth0 is the “heartbeat” interface for the cluster. Use the eth0 interface on each OS-E system as the connection to the cluster.

```
NNOS-E> config cluster
config cluster> config box 1
config box 1> set identifier 00:04:23:d7:9f:34
config box 1> config interface eth0
config interface eth0> config ip heartbeat
Creating ‘ip heartbeat’
config ip heartbeat> set ip-address static 192.168.1.1/24
config ip heartbeat> config telnet
config telnet> return
config ip heartbeat> config ssh
```

```
config ssh> return
config ip heartbeat> config bootp-server
config bootp-server> return
config ip heartbeat> config vrrp
config vrrp> return
```

Note: Optionally, you can run the **config setup** script to configure the IP addresses, management port, and other settings presented in the script.

By configuring messaging on the OS-E master, the master looks through the configurations of all other devices to find out which interface is used for messaging. (If multiple interfaces are configured, the master only communicates with one—the first it finds.) The master then communicates with the identified interface to share configuration and data.

```
config ip heartbeat> config messaging
config messaging> set admin enabled
config messaging> set certificate vsp tls
  certificate 208.45.178.216.pfx
config messaging> set port 5312
config messaging> set protocol tls
config messaging> return
config ip heartbeat> return
config interface eth0> return
config box 1>
```

Configure the interface and the protocols over which you will run management sessions to the OS-E. This is an “out-of-band” interface that allows you to separate management traffic from SIP signaling and media streams.

```
config box 1> config interface eth3
Creating ‘interface eth3’
config interface eth3> config ip mgmt
Creating ‘ip mgmt’
config ip mgmt> set ip address static 192.168.2.1/24
config ip mgmt> config ssh
config ssh> return
config ip mgmt> config web
config web> set protocol https 443 0
config web> return
config ip mgmt> config sip
config sip> set udp-port 5060
config sip> return
config ip mgmt> config icmp
config icmp> return
config ip mgmt> config media-ports
config media-ports> return
config ip-mgmt> return
config interface eth3> return
```

```
config box 1> config cli
config cli> set prompt nn2610-1
config cli> set banner ""
config cli> set display paged 50
config cli> return
config box 1> return
config cluster>
```

5. Configure the second OS-E system in the cluster. Note that you also configure eth0 as the "heartbeat" interface to the cluster.

```
   config cluster> config box 2
config box 2> set hostname nn2610-2
config box 2> set name ""
config box 2> set contact ""
config box 2> set location ""
config box 2> set identifier 00:04:23:c3:22:f4
config box 2> config interface eth0
config interface eth0> config ip heartbeat
config ip heartbeat> set ip-address static 192.168.1.2/24
config ip heartbeat> config telnet
config telnet> return
config ip heartbeat> config ssh
config ssh> return
config ip heartbeat> config web
config web> set protocol https 443 0
config web> return
config ip heartbeat> config icmp
config icmp> return
config ip heartbeat> config vrrp
config vrrp> return
config ip heartbeat> config messaging
config messaging> set admin enabled
config messaging> set certificate vsp tls
  certificate 208.45.178.216.pfx
config messaging> set port 5312
config messaging> set protocol tls
config messaging> return
config ip heartbeat> return
config interface eth0> return
config box 2>
```

Configure the interface and the protocols over which you will run management sessions. This is an "out-of-band" interface that allows you to separate management traffic from SIP signaling and media streams.

```
config box 2> config interface eth3
Creating 'interface eth3'
config interface eth3> config ip mgmt
Creating 'ip mgmt'
config ip mgmt> set ip address static 192.168.2.2/24
```

```
config ip mgmt> config ssh
config ssh> return
config ip mgmt> config web
config web> set protocol https 443 0
config web> return
config ip mgmt> config sip
config sip> set udp-port 5060
config sip> set nat-translation enabled
config sip> return
config ip mgmt> config icmp
config icmp> return
config ip mgmt> config media-ports
config media-ports> return
config ip-mgmt> return
config interface eth3> return
config box 1> config cli
config cli> set prompt NNOS-E-2
config cli> set banner ""
config cli> set display paged 50
config cli> return
config box 1> return
config cluster> set share media-ports true
config cluster> set share signaling-entries true
config cluster> set mirror-media-streams true
```

6. Configure VRRP on the OS-E interfaces to handle the public and private sides of the network. Note that the first VRRP interface connects the public side; the second VRRP interface connects the private side.

A VRRP configuration for IP interfaces includes a list of box/interface pairings. The first pair in this list is the *primary interface*. The second pair in the list is the *backup interface* and will take over if the primary goes down. You can configure additional levels of redundancy by specifying more box/interface pairs of lower priority. Priority is based on the positioning of the **set host-interface** command.

```
config cluster> config vrrp
config vrrp> config vinterface vx0
config vinterface vx0> set group 1
...vinterface vx0> set host-interface cluster box 1 interface eth1
...vinterface vx0> set host-interface cluster box 2 interface eth1
config vinterface vx0> config ip public
Creating 'ip public'
config ip public> set ip-address static 208.45.178.216/28
config ip public> config ssh
config ssh> return
config ip public> config web
config web> set protocol https 443 0
config web> return
config ip public> config sip
config sip> set nat-translation enabled
```

```
config sip> set udp-port 5060
config sip> set tcp-port 5060
config sip> set tls-port 5061
config sip> set certificate vsp\tls\certificate 208.45.178.216.pfx
config sip> return
config ip public> config icmp
config icmp> return
config ip public> config media-ports
config media-ports> return
config ip public> config routing
config routing> config route default
Creating 'route default'
config route default> set gateway 208.45.178.209
config route default> return
config routing> return
config ip public> return
config vinterface vx0> return
config vrrp>

config cluster> config vrrp
config vrrp> config vinterface vx1
config vinterface vx1> set group 1
...vinterface vx1> set host-interface cluster box 1 interface eth2
...vinterface vx1> set host-interface cluster box 2 interface eth2
config vinterface vx1> config ip private
Creating 'ip private'
config ip private> set ip-address static 208.45.178.216/28
config ip private> config ssh
config ssh> return
config ip public> config web
config web> set protocol https 443 0
config web> return
config ip private> config sip
config sip> set nat-translation enabled
config sip> set udp-port 5060
config sip> set tcp-port 5060
config sip> set tls-port 5061
config sip> set certificate vsp\tls\certificate 208.45.178.216.pfx
config sip> return
config ip private> config icmp
config icmp> return
config ip private> config media-ports
config media-ports> return
config ip private> config routing
config routing> config route static-to-asx
Creating 'route static-to-asx'
config route static-to-asx> set destination network 208.45.178.0/24
config route static-to-asx> set gateway 10.144.10.254
config route static-to-asx> return
config routing> return
config ip private> return
```

```
config vinterface vx1> return
config vrrp> return
config cluster> return
```

7. Configure the master-services that you want to run on the cluster.

```
config> config master-services
config master-services> config accounting
config accounting> set host-box cluster\box 1
config accounting> set host-box cluster\box 2
config accounting> set group 1
config accounting> return
config master-services> config database
config database> set host-box cluster\box 1
config database> set host-box cluster\box 2
config database> set group 1
config database> set media enabled
config database> return
config master-services> return
config>
```

8. For TLS, you will need to upload the TLS certificate file on each OS-E system in the cluster. Copy the certificate that you receive from the CA to the OS-E using a compatible file transfer mechanism, such as PuTTY Secure Copy (PSCP). If you have the file on a local network PC, use PSCP to move the file to a directory path on the OS-E.

   The following example PSCP command copies the certificate file named **208.45.178.216.pfx** from the PC root directory to the OS-E system at IP address **208.178.216.pfx** in the directory **/cxc/certs/208.45.178.216.pfx**.

```
C:\ pscp -l root -pw sips -P 2200 208.45.178.216.pfx 208.45.178.216:/
   cxc/certs/208.45.178.216.pfx
```

   The following CLI session sets the directory and certificate file name path, specifies the passphrase, and whether to allow SSL Version 2 operability.

```
NNOS-E> config vsp
config vsp> config tls
config tls> config certificate 208.45.178.216.pfx
config certificate 208.45.178.216.pfx> set allow-sslv2 true
config certificate 208.45.178.216.pfx> set certificate-file /cxc/
   certs/208.45.178.216.pfx.pfx
config certificate 208.45.178.216.pfx> set passphrase-tag pass
```

   By default, the OS-E only supports SSLv3 or TLSv1. If you require SSLv2 for interoperability, set this property **true**. Specify the passphrase-tag associated with the certificate file. Use this property if the certificate file is encrypted to have its private key information protected. This passphrase tag must match the string with which the certificate was encrypted.

9. Power up the other OS-E systems in the cluster and connect them to the network. This initiates a configuration download from the cluster master so the systems acquire their initial configuration (IP addresses, etc.).
10. Use the CLI or OS-E Management System at the cluster master to configure any additional features. These features include the objects and settings under the VSP object, including:
    - default-session-config
    - registration-plan
    - dial-plan
    - enterprise servers, carriers, and gateways

Configuring External Messaging

Messaging is the mechanism the OS-E uses to communicate among boxes in a cluster. Messaging sets up a listening socket on an interface, enabling the interface to receive messaging traffic and participate in clustering and media partnering.

In a cluster, the master looks through the configurations of all OS-E systems to find out which interface is used for messaging. (If multiple interfaces are configured, the master only communicates with one—the first it finds.) The master then communicates with the identified interface to share configuration and data.

In media partnering, you configure a specific IP address (on a different box) as a partner. On the box that owns that IP address, you need to configure and enable messaging for media partnering to operate.

CLI Session

The following CLI session configures messaging on box 1, interface eth0.

```
NNOS-E> config cluster
config cluster> config box 1
config box 1> config interface eth0
config interface eth0> config ip boston1
config ip boston1> config messaging
config messaging> set admin enabled
config messaging> set certificate vsp tls certificate name
config messaging> set port 13002
config messaging> set protocol tls
```

Configuring Cluster Load Balancing

Load balancing of SIP processing across cluster interfaces requires both headend and backing interfaces. The *headend* interface is the central distribution point. It does not do any SIP processing, it only forwards the calls to its configured backing interfaces. When you configure a SIP phone, you would configure it to point to the headend interface.

To configure an IP interface as a headend interface, you simply configure the **sip** object with backing interfaces. Their presence contained within the IP configuration results in the interface being treated by the OS-E as a headend interface.

The *backing-interfaces* are identified as such within this **sip** object. In the **backing-interface** property, you reference previously configured IP interfaces. The backing interface is the location at which the OS-E terminates TCP and TLS connections (and where UDP transport messages arrive) and handles SIP processing. The OS-E uses round-robin load-balancing to distribute message across the configured backing interfaces.

To correctly configure load-balancing for SIP processing, you must do the following:

1. Configure the IP interfaces that will be used for both the headend and backing interfaces.
2. The SIP properties of the backing interfaces must match those of the head interface. For example, they must all use the same port assignments, and if you are using TLS, they must all use the same certificate.
3. You must enable the **master-services registration** object so that the interfaces can share the registration database.

To verify your configuration, first ensure that all SIP properties match. From the CLI at the headend, execute the **show load-balance** command. This lists all associated backing interfaces (and statistics). From each box hosting a backing interface, execute **show backing-interface** to display configuration and statistics information.

The following CLI session assumes that you have configured a three-box cluster, with box 1 containing the headend interface, with boxes 2 and 3 containing the backing interfaces over which traffic is load balanced. This session sets the backing interfaces for load balancing SIP traffic that is distributed from the headend interface at IP address 215.2.3.0/24.

CLI Session

```
config> config cluster
config cluster> config box 1
config box 1> config interface eth1
config interface eth1> config ip public
Creating 'ip public'
config ip public? set ip-address static 215.2.3.0/24
config ip public> config sip
config sip> config load-balancing
config load-balancing> set backing-interface cluster box 2 interface
   eth1 ip public
Creating 'cluster\box 2\interface eth1\ip public'
config load-balancing> set backing-interface cluster box 3 interface
   eth1 ip public

config sip> show
cluster
 box 1
  interface eth1
   ip public
    sip
     admin enabled
     backing-interface cluster\box 2\interface eth1\ip public2
     backing-interface cluster\box 3\interface eth1\ip public3


NNOS-E> show load-balance

Head-end IP 215.2.3.0:  undersubscribed:
-------------------------------------------------------------------------------
Backing IP      State       Added   Removed   Maximum   Current  Percent
-------------------------------------------------------------------------------
215.6.7.0       Down            0         0         0         0     0.0%
215.8.9.0       Down            0         0         0         0     0.0%
-------------------------------------------------------------------------------
Totals:                         0         0         0         0   100.0%

NNOS-E>
```

Restarting an OS-E Cluster

You can perform a simultaneous warm restart of all systems in a cluster by using the **restart cluster** command. A warm restart simply restarts the OS-E applications on each system without rebooting the operating system.

If you warm restart an individual device in the cluster, the OS-E automatically rejoins the cluster when it comes back up. If that box is hosting a master service or a VRRP interface, the service or interface may fail over to a different OS-E system.

If you need to shut a system down by turning the power off, use the **restart halt** command before pressing the power button or disconnecting the power source. A **restart halt** will properly prepare a system for a shutdown. The OS-E system will rejoin the cluster when it comes back up.

Chapter 5. Installing Certificates and Commissioning TLS Networks

About This Chapter

This chapter provides information on commissioning the OS-E to run the Transport Layer Security protocol (TLS) over Ethernet interfaces.

TLS Overview

TLS is an encapsulation (and cryptographic) protocol that provides privacy and security between communicating applications over the Internet. The OS-E uses TLS to authenticate SIP users and to encrypt/decrypt SIP traffic across participating carrier and enterprise SIP applications.

For a complete description of the TLS protocol, refer to the following RFCs:

- RFC 2246, The TLS Protocol Version 1.0
- RFC 3261, Session Initiation Protocol (see Section 26.3.1)

The image below illustrates a sample network running TLS on Ethernet interfaces to both the private network and the public Internet.

Steps to Configuring TLS

To configure the private and public network interfaces in the image below, with TLS you need to perform the following steps:

1. If not already done, configure the management interfaces, network routes, protocols, and services using the **cluster/box**, **master-services**, and **services** objects.
2. If not already done, install a signed certificate from a valid Certificate Authority (CA). Go the section, “Before Configuring TLS.”.
3. Configure the certificate using the **tls/vsp** configuration object. Go to the section, “Configuring the Certificate on the OS-E.”
4. Configure the SIP protocol on each interface to use the installed certificate. Go to the section, “Configuring TLS on Ethernet Interfaces.”

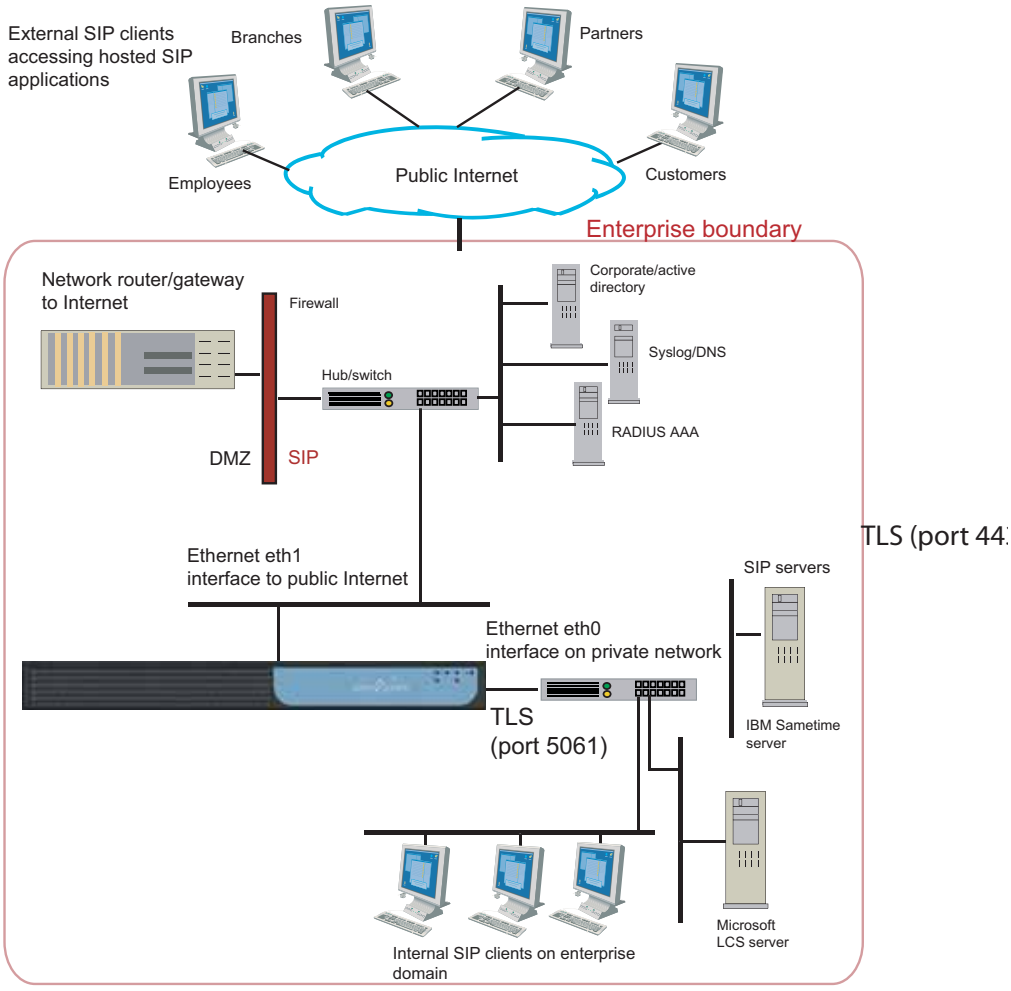

Note: If you are operating with Microsoft Live Communications Server (LCS), refer to the *Net-Net OS-E – System Administration Guide* for information on installing, importing, and exporting certificates.

Before Configuring TLS

You will need to install the required X.509 certificate(s) for the TLS protocol and SIP session establishment. A certificate includes the *.cer* certificate file name and the encrypted keys, mathematically related private and public data keys indexed by a unique name. A private key is kept secure—never displayed and never transmitted over the network. A public key, when bound to a fully qualified domain name (FQDN) by an authorized certification authority (CA), becomes an X.509 certificate.

If you do not have a certificate and encrypted key for your network, use this section to create and install a certificate. See the following table below for a summary of required steps.

Task	See this section	CLI command
1. Using the OS-E software, create a self-signed X.509 certificate and encrypted key.	*Creating a self-signed certificate and key pair from the OS-E*	**`cert-gen`**
2. Using the self-signed certificate that you created in Step 1, generate a Certification Signing Request (CSR) in PEM format.	*Generating a Certification Signing Request (CSR)*	**`cert-request`**

Task	See this section	CLI command
3. Sign the CSR using one of these two methods: • Use a a valid CA, like VeriSign (required if a “trusted” certificate is necessary) to sign the CSR **or** • Use OpenSSL to sign the CSR	*Signing a CSR using either a valid CA or OpenSSL*	N/A
4. When you receive the signed certificate, use the OS-E software to load the signed certificate onto the system.	*Updating the self-signed certificate*	**`cert-update`**

Step 1. Creating a Self-Signed Certificate and Key Pair from the OS-E

Use the OS-E software to generate a cryptographic key pair and a self-signed X.509 certificate in PKCS#12 format. Once you create a self-signed certificate, you can generate the Certification Signing Request, a portion of which will be required by the CA upon submission of their form.

Under **Actions**, select **cert-gen** from list. The following image illustrates the OS-E Management System Generate new key and certificate page.

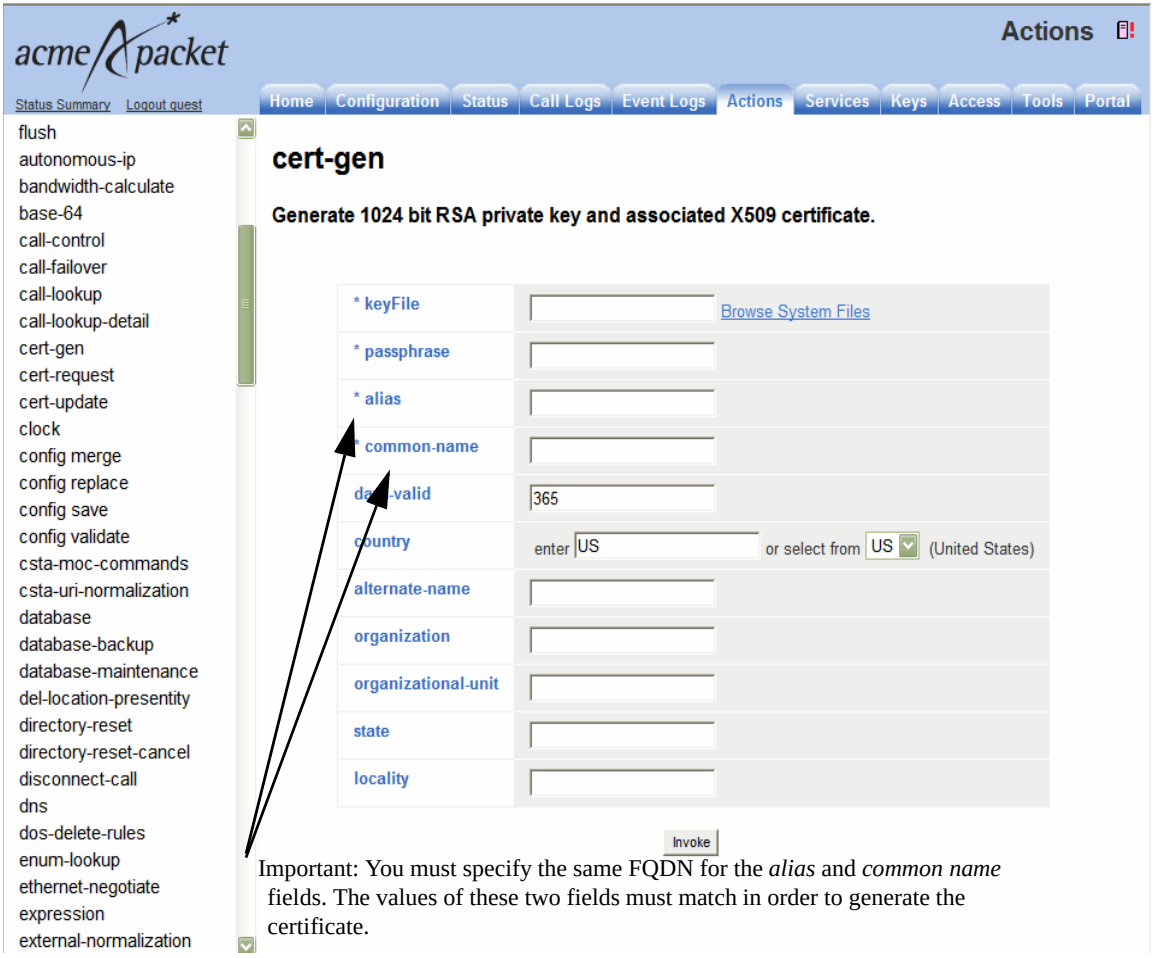

Important: You must specify the same FQDN for the *alias* and *common name* fields. The values of these two fields must match in order to generate the certificate.

Complete the fields on the **Generate new key and certificates** page, as follows:

keyFile—Specify the name and directory path of the resulting key name that you want to use, along with the p12 or .pfx file extension.

Example: /cxc/certs/<*myNetworkKey*>.p12

passphrase—Specify a password to be associated with the self-signed certificate. The text that you specify will be encrypted in the certificate.

alias—Specify the FQDN of the OS-E system using this certificate, such as *nn2610.acmepacket.com*. Omit HTTP:// and HTTPS://. This allows the certificate to be referenced.

Note: The value (FQDN) you enter for the **alias** field must be identical to the value you enter for the **common-name** field.

common-name—Specify the FQDN of the OS-E system using this certificate, such as *nn2610.acmepacket.com*. Omit HTTP:// and HTTPS://. Do not use your personal name in this field. The common name is a component attribute of the certificate's *distinguished name*.

days-valid—Enter the number of days for which the certificate is valid. If your certificate if effective for one year, then enter the number 365

country—Select the ISO country code: US (United States), AU (Australia), IN (India), IT (Italy), UK (United Kingdom), CA (Canada). The country is a component attribute of the certificate's *distinguished name*.

alternate-name—Optional; this usually a name that complies with the ASN.1 specification, such as a DNS name, IP address, URI, etc.

organization—Optional. Enter the name under which your business is legally registered. The listed organization must be the legal registrant of the domain name in the certificate request. If you are enrolling as a small business/sole proprietor, enter the certificate requestor's name in the "Organization" field, and the DBA (doing business as) name in the "Organizational Unit" field. The organization is a component attribute of the certificate's *distinguished name*.

organizational-unit—Optional. Use this field to differentiate between divisions within an organization. For example, "Engineering" or "Human Resources." If applicable, you may enter the DBA (doing business as) name in this field. The organizational unit is a component attribute of the certificate's *distinguished name*.

state—Optional; if in the US, enter one of the fifty state names in full where your organization is located, such as Massachusetts; if outside the US, enter the full name of a province or region.

locality—Optional; enter the name of a city.

When you are finished filling out the fields, click **Invoke**. A Success message appears.

Viewing the Certificate

To view the self-signed certificate, select the **Keys** tab from the main menu bar, then select the keyFile that you just created from the Key Stores list on the left. Click View to display the Certificate Properties page, as shown in the following image.

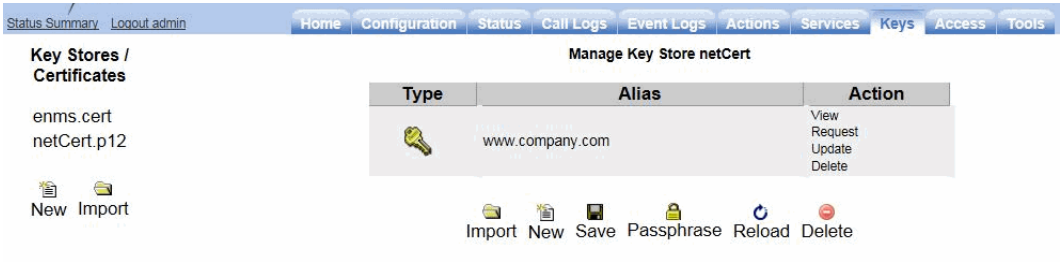

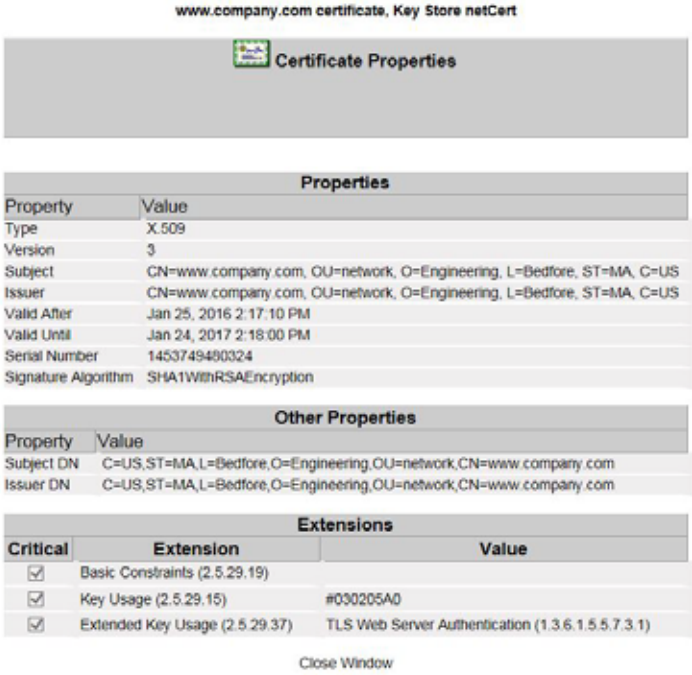

Step 2. Generating a Certification Signing Request

After you create the self-signed certificate from Step 1, you must generate a certification signing request (CSR) that you can submit to the CA for the X.509 certificate. Select the **request** action.

The Generate Certificate Signing Request page and the resulting certificate signing request appear. Enter the password that you created in Step 1 in the **passphrase** field and click on **Generate Certificate Signing Request**. Select the key store file on the left and enter your password when prompted. Manage Key Store is displayed; click on **Request** under **Action**.

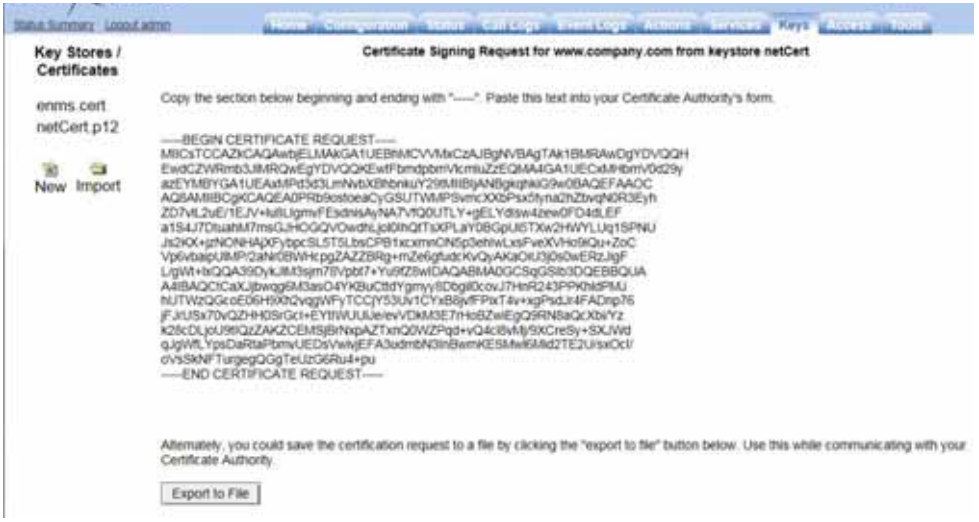

Click the **Export to File** button to save the CSR provided by the CA to external file.

If you choose to create a CSR in a PEM-formatted file, select the **cert-request** action. The file contains the same request as shown in the following image.

Complete the fields on the **Generate Certification Request** page, using the same settings that you invoked from Step 1, as follows:

key-file—Specify the name and OS-E directory path of the resulting key name that you want to use, along with the p12 or .pfx file extension. This is a mandatory field. Example: /cxc/certs/myNetworkKey.p12

passphrase—Specify a password to be associated with the certificate issued by the CA. The text that you specify will be encrypted in the CSR.

alias—Optional. However, the value you enter for the **alias** field must be identical to the value you enter for the **common-name** field.

csr-file—Specify the name and directory path of the resulting CSR file. This is the file from which you will cut and paste the required information for the CA at the time that you submit the certificate request. By default, the CSR file resides in the directory named /cxc/certs.

When you are finished filling out the fields, click **Invoke**. A Success message appears.

Viewing the .CSR File

Since the .cer file is in PEM format, you can open the file using a text editor.

Step 3. Signing a CSR Using Either a Valid CA or OpenSSL

After you generate the CSR, you must sign the CSR using *one* of two methods. You can either:

— Sign the CSR using a well-known CA, for example, VeriSign.

or

— Sign the CSR using OpenSSL.

This section describes how to sign the CSR using either method.

Note: If your network requires a “trusted” certificate, then follow the instructions below to sign the CSR using a valid, well-known CA.

Using a Certification Authority to Sign the CSR

You get the signed X.509 certificate from a valid CA, such as VeriSign. The CA issues a certificate stating and guaranteeing that the key contained in the certificate belongs to the person or organization noted in the certificate. The CA verifies the identify of the applicant's so that users can trust certificates issued by that CA to belong to the people and data identified in it, and not to an imposter.

Certificate Formats

The OS-E certificate file can be in the following formats:

- PKCS#12—Public Key Cryptography Standard #12 format from Microsoft IIS Version 5 (binary)
- PEM—Privacy-enhanced mail (PEM) encoded format from any OpenSSL-based Web server (ASCII)

Using OpenSSL to Sign the CSR

This section provides information on how you can generate a self-signed certificate for testing TLS with the OS-E using OpenSSL. This is an alternative method to using a valid CA to sign the CSR.

This section describes how to do the following things:

- Create an OpenSSL Certificate Authority (CA).
- Generate a private key and CSR on the OS-E system and sign in with the OpenSSL CA.
- Generate a Private Key and CSR without the OS-E system (not supported).

- Use OpenSSL to convert an X.509 certificate and/or RSA key to a Public-Key Cryptography Standard #12 (PKCS#12) format.

Note: Before using this method, download the OpenSSL program and install it on a Unix/Linux or Windows system. You also need to add the location of the OpenSSL executables to the PATH. In a Windows environment, this will need to do this manually, requiring a reboot to take effect.

Creating an OpenSSL Certificate Authority (CA)

To create and Open SSL Certificate Authority (CA) on a Unix/Linux system, perform all steps as "root." On a Windows system, perform all steps as "Administrator."

1. **Create directories to store certificates.**

   The main CA folder is the directory where the **C**ertificate **A**uthority files will reside. The "private" directory stores the private keys. The "certs" directory stores the certificates (or public keys). The "csrs" directory stores the Certificate Signing Requests.

On Unix:

```
mkdir /CA
mkdir /CA/private
mkdir /CA/csrs
mkdir /CA/certs
```

Windows (cmd):

```
mkdir C:\CA
mkdir C:\CA\private
mkdir C:\CA\csrs
mkdir C:\CA\certs
```

2. **Create files to support the generation process.**

   Create the "index.txt" file with no contents. This is the database to which OpenSSL keeps track of generated certificates generated. Create the "serial" file with a number so that each generated certificate is labeled with a number for tracking purposes.

Unix:

```
touch /CA/index.txt
```

```
echo 01 > /CA/serial
```

Windows (cmd):

```
copy con C:\CA\index.txt
echo 01 > C:\CA\serial
```

3. **Create the OpenSSL configuration file.**

Note: The following are example values only.

Unix:

Using a text editor, create "/CA/openssl.cnf."

```
[ ca ]
default_ca      = local_ca

[ local_ca ]
dir             = /CA
certificate     = $dir/certs/ca.cer
database        = $dir/index.txt
new_certs_dir   = $dir/certs
private_key     = $dir/private/ca.key
serial          = $dir/serial



default_crl_days        = 365
default_days            = 365
default_md              = md5

policy          = local_ca_policy
x509_extensions = local_ca_extensions

[ local_ca_policy ]
commonName              = supplied
stateOrProvinceName     = optional
countryName             = optional
emailAddress            = optional
organizationName        = optional
organizationalUnitName  = optional

[ local_ca_extensions ]
basicConstraints        = CA:true
nsCertType              = server

[ root_ca_extensions ]
basicConstraints        = CA:true
nsCertType              = server
```

```
[ req ]
default_bits    = 2048
default_keyfile = /CA/private/ca.key
default_md      = md5

prompt                  = yes
distinguished_name      = root_ca_distinguished_name
x509_extensions         = root_ca_extensions

[ root_ca_distinguished_name ]
countryName                     = Country Name (2 letter code)
countryName_default             = US
countryName_min                 = 2
countryName_max                 = 2

stateOrProvinceName             = State or Province Name (full name)
stateOrProvinceName_default     = MA

localityName                    = Locality Name (eg, city)
localityName_default            = Maynard

0.organizationName              = Organization Name (eg, company)
0.organizationName_default      = Acme Packet, Inc.

organizationalUnitName          = Organizational Unit Name
  (eg,section)
organizationalUnitName_default  = Support

commonName                      = Common Name (eg, YOUR name)
commonName_max                  = 64

emailAddress                    = Email Address
emailAddress_default            = jgentile@acmepacket.com
emailAddress_max                = 64

[ req_attributes ]
challengePassword               = A challenge password
challengePassword_min           = 4
challengePassword_max           = 20

unstructuredName                = An optional company name
```

Windows:

Using a text editor, create “C:\CA\openssl.cnf.”

Note: The following are example values only.

```
[ ca ]
default_ca      = local_ca

[ local_ca ]
dir             = C:\\CA
certificate     = $dir\\certs\\ca.cer
database        = $dir\\index.txt
new_certs_dir   = $dir\\certs
private_key     = $dir\\private\ca.key
serial          = $dir\\serial

default_crl_days        = 365
default_days            = 365
default_md              = md5

policy          = local_ca_policy
x509_extensions = local_ca_extensions

[ local_ca_policy ]
commonName              = supplied
stateOrProvinceName     = optional
countryName             = optional
emailAddress            = optional
organizationName        = optional
organizationalUnitName  = optional

[ local_ca_extensions ]
basicConstraints        = CA:false
nsCertType              = server

[ root_ca_extensions ]
basicConstraints        = CA:true
nsCertType              = server

[ req ]
default_bits    = 2048
default_keyfile = C:\\CA\\private\\ca.key
default_md      = md5

prompt                  = yes
distinguished_name      = root_ca_distinguished_name
x509_extensions         = root_ca_extensions

[ root_ca_distinguished_name ]
countryName                     = Country Name (2 letter code)
countryName_default             = US
countryName_min                 = 2
countryName_max                 = 2

stateOrProvinceName             = State or Province Name (full name)
stateOrProvinceName_default     = MA
```

```
localityName                    = Locality Name (eg, city)
localityName_default            = Maynard

0.organizationName              = Organization Name (eg, company)
0.organizationName_default      = Acme Packet, Inc.

organizationalUnitName          = Organizational Unit Name (eg,
                                section)
organizationalUnitName_default  = Support

commonName                      = Common Name (eg, YOUR name)
commonName_max                  = 64

emailAddress                    = Email Address
emailAddress_default            = jgentile@acmepacket.com
emailAddress_max                = 64

[ req_attributes ]
challengePassword               = A challenge password
challengePassword_min           = 4
challengePassword_max           = 20

unstructuredName                = An optional company name
```

4. **Generate the CA's private key and Master Certificate (public key).**

   This step will generate two files:

   - **CA/private/ca.key** (**C:\CA\private\ca.key** on Windows) – This is the CA's private key used to sign certificates. Keep this secure. If this key is compromised, it can be used to create certificates for malicious purposes.
   - **CA/certs/ca.cer** (**C:\CA\certs\ca.cer** on Windows) – This is the CA's certificate (public key). This is the file that would be distributed to client's "Trusted Root" stores to trust any certificates signed by this CA's private key.

Unix:

```
openssl req -x509 -new -config /CA/openssl.cnf –days 3000 -out /CA/
certs/ca.cer
```

Windows (cmd):

```
openssl req -x509 -new -config C:\CA\openssl.cnf –days 3000 -out
   C:\CA\certs\ca.cer
```

The **ca.key** is created automatically based on the configuration file.

Enter a strong passphrase for the CA key. Remember it, as this helps protect the security of your CA.

Fill in the following fields:

```
Country Name (2 letter code) [US]: <your country>
State or Province Name (full name) [MA]: <your state/province>
Locality Name (eg, city) [Maynard]: <your locale>
Organization Name (eg, company) [Acme Packet]: <your company>
Organizational Unit Name (eg, section) [Support]: <your department>
Common Name (eg, YOUR name) []: <Use the FQDN of the CA system
  running OpenSSL>
Email Address []: <your email address>
```

Note: The "Common Name" field is the most important. This is the name that will be provided to the CA, so use the Fully Qualified Domain Name (FQDN) of the system on which you are running OpenSSL.

5. **Change permissions on the CA's key to only allow "root" access**:

Unix:

```
chmod 700 /CA/private/ca.key
```

Windows (cmd):

```
echo y|cacls C:\CA\private\ca.key /G %COMPUTERNAME%\Administrator:F
```

Note: You should only need to complete the process for setting up the CA once, while the processes for signing Certificates must be repeated every time a certificate needs to be generated.

Generating a Private Key and Certificate Signing Request (CSR) with the OS-E

To generate a private key and CSR on the OS-E and sign in with the OpenSSL CA, perform the following steps:

1. **Create a CSR on the OS-E.**

   Refer to the section in the chapter, "Before Configuring TLS.":

   - Use the **cert-gen** utility to generate a Self-Signed Certificate (known as a private key) in PKCS#12 format. In this example, the file name is *cxc.pfx***.**

- Use the **cert-request** utility to generate a Certificate Signing Request (CSR) in PEM format on the appropriate OS-E system. In this example, the file name is *cxc.csr.*

Note: The "common-name" field on the "cert-gen" page is the most important. This is the name that will be used to validate the certificate. Use the Fully Qualified Domain Name (FQDN) of the appropriate OS-E system, such as nnose.acmepacket.com.

Currently, some phones, such as Eyebeam do not support wildcard certificates where the common-name uses an asterisk character (*) in the domain name, such as *.acmepacket.com.

These files are created in the **/cxc/certs/** directory on the OS-E.

2. **Copy the CSR to the OS-E.**

   Download the .csr file generated on the OS-E, and then copy it to the CA system into the **/CA/csrs/** directory. For a Windows system, copy it to the **C:\CA\csrs** directory.

3. **Sign the CSR with your OpenSSL CA.**

Unix

```
openssl ca -config /CA/openssl.cnf -in /CA/csrs/cxc.csr –out /CA/
certs/cxc.pem
```

Windows

```
openssl ca -config C:\CA\openssl.cnf -in C:\CA\csrs\cxc.csr –out
C:\CA\certs\cxc.pem
```

Enter the pass phrase for the CA key

Respond "y" to the questions to generate and commit.

4. **Update the private key (cxc.pfx) with the signed public key (cxc.pem) on the system.**

   Upload the newly generated *cxc.pem* file back to the OS-E , as covered earlier in this chapter. Refer to the section in the chapter, "Before Configuring TLS."

   - Use the **cert-update** utility to update the "/cxc/certs/cxc.pfx" file on the OS-E with the "/cxc/certs/cxc.pem" file.

- Configure a TLS certificate, as covered earlier in this chapter. Be sure to associate it with the SIP protocol on the appropriate network interface.

Note: You can use the /CA/certs/ca.cer (C:\CA\certs\ca.cer on Windows) file to import into a "Trusted Root Store." For example, you can install this in Windows (Internet Explorer) for use with Soft Phones, such as Eyebeam. If you deploy the ca.cer file to multiple systems into the "Trusted Root Store", then those systems will "trust" any certificates signed by this CA.

Generating a Private Key and Certificate Signing Request (CSR) without the OS-E

Instead of generating the private key and CSR on the OS-E, you can generate it using OpenSSL exclusively. This is not the supported method.

1. **Create a CSR and Private Key for the OS-E System**

Unix:

```
openssl req -new -config /CA/openssl.cnf -out /CA/csrs/cxc_csr.pem
-keyout /CA/certs/cxc_pk.pem
```

Windows (cmd:)

```
openssl req -new -config C:\CA\openssl.cnf -out
C:\CA\csrs\cxc_csr.pem -keyout C:\CA\certs\cxc_pk.pem
```

Use the OpenSSL "req" utility to generate a Self-Signed Certificate (private key) and the Certificate Signing Request (CSR) in PEM format. In this example, the file names are *cxc_pk.pem* for the private key, and *cxc_csr.pem* for the CSR.

Enter a pass phrase for the CA key, and complete he following fields::

```
Country Name (2 letter code) [US]: <your country>
State or Province Name (full name) [MA]: <your state/province>
Locality Name (eg, city) [Maynard]: <your locale>
Organization Name (eg, company) [Acme Packet]: <your company>
Organizational Unit Name (eg, section) [Support]: <your department>
Common Name (eg, YOUR name) []: <Use the FQDN of the CXC>
Email Address []: <your email address>
```

Note: The "common-name" field is the most important entry. This is the name that will be used to validate the certificate. Use the Fully Qualified Domain Name (FQDN) of the appropriate OS-E system, such as nnose.acmepacket.com.

Currently, some phones, such as Eyebeam do not support wildcard certificates where the common-name uses an asterisk character (*) in the domain name, such as *.acmepacket.com

2. **Sign the CSR with your OpenSSL CA.**

Unix

```
openssl ca -config /CA/openssl.cnf -in /CA/csrs/cxc_csr.pem –out /
CA/certs/cxc.pem
```

Windows

```
openssl ca -config C:\CA\openssl.cnf -in C:\CA\csrs\cxc_csr.pem
–out C:\CA\certs\cxc.pem
```

Enter the pass phrase for the CA key, then respond “y” to the questions to generate and commit.

3. **Merge the Private Key and Signed Public Key into one file**.

Unix

```
cat /CA/certs/cxc.pem /CA/certs/cxc_pk.pem > /CA/certs/cxc.list.pem
```

Windows (cmd)

```
copy /CA/certs/cxc.pem + /CA/certs/cxc_pk.pem /CA/certs/
cxc.list.pem
```

4. Upload the newly generated *cxc.list.pem* file back to the OS-E, then configure a TLS certificate, as covered earlier in this chapter. Be sure to associate it with the SIP protocol on the appropriate network interface.

Note: You can use the /CA/certs/ca.cer (C:\CA\certs\ca.cer on Windows) file to import into a “Trusted Root Store.” For example, you can install this in Windows (Internet Explorer) for use with Soft Phones, such as Eyebeam. If you deploy the ca.cer file to multiple systems into the “Trusted Root Store”, then those systems will “trust” any certificates signed by this CA.

Using OpenSSL to Convert X.509 and RSA Keys

This section describes how to use OpenSSL to convert an X.509 certificate and/or RSA key to a Public-Key Cryptography Standard #12 (PKCS#12) format.

Requirements

You must have a working installation of the OpenSSL software and be able to execute OpenSSL from the command line.

Refer to "CTX106627 - How to Install the OpenSSL Toolkit," for more information on obtaining and installing OpenSSL.

The PKCS#12 specifies a portable format for storing and transporting certificates, private keys, and miscellaneous secrets. It is the preferred format for many certificate handling operations and is supported by most browsers and recent releases of the Windows family of operating systems. It has the advantage of being able to store the certificate and corresponding key, root certificate, and any other certificates in the chain in a single file.

Procedure

1. Ensure that the certificate(s) and key are in PEM format.

   - **To convert a certificate from DER to PEM:**

```
x509 –in input.crt –inform DER –out output.crt –outform PEM
```

   - **To convert a key from DER to PEM:**

```
rsa –in input.key –inform DER –out output.key –outform PEM
```

   - **To convert a key from NET to PEM:**

```
rsa –in input.key –inform NET –out output.key –outform PEM
```

Note: The obsolete NET (Netscape server) format is encrypted using an unsalted RC4 symmetric cipher so a passphrase will be requested. If you do not have access to this passphrase it is unlikely you will be able to recover the key

2. Use the **openssl** command to read the PEM encoded certificate(s) and key and export to a single PKCS#12 file as follows:

```
openssl pkcs12 -export -in input.crt -inkey input.key -out
   bundle.p12
```

Note: By default, the key will be encrypted with Triple DES so you will be prompted for an export password (which may be blank).

The PEM formatted root certificate and any other certificates in the chain can be merged into a single file such as root.crt, and included in the PKCS#12 file as follows:

```
openssl pkcs12 -export -in input.crt -inkey input.key -certfile
root.crt -out bundle.p12
```

Step 4. Updating the Self-Signed Certificate

The **cert-update** action allows you to load the signed certificate that you receive from the CA. Once you have received the file, perform the following steps:

1. Upload the file to the OS-E using the **Tools**/**Upload file** function to browse for CA's certificate. Specify the destination path on the OS-E system, such as /cxc/certs, and specify the destination name of the certificate.
2. Select the **Keys** tab and select the appropriate key from the Key Stores list to display the Manage Key Store page.
3. Click **Update** to browse for the file that you uploaded in Step 1.
4. Click **Update** to load the signed certificate to the CXC.

If you choose to update the certificate using the **cert-update** action rather than from the **Keys** tab, complete the fields as follows:

keyFile—Specify the name and directory path of the key that you want to update.

Example: /cxc/certs/<*myNetworkKey*>.p12

alias—Optional. Specify the alias for the keyFile name, if previously created.

password—Specify the password associated with the keyFile, as specified previously.

certFile—Specify the name and directory path of the signed certificate that you received from the CA and uploaded to the OS-E using the OS-E Management System **Tools/Upload File** function or other file transfer mechanism.

Subject Alternative Name for HTTPS Certificates Support

The OS-E supports Subject Alternative Name (SAN) for use with HTTPS certificates. SAN is a X509 version 3 certificate extension that allows one to specify a list of host names protected by a single SSL certificate.

To add multiple SANs to a certificate:

1. Select the **Keys** tab and either click **New** to create a new key store or select the existing store on which you want to add a certificate.

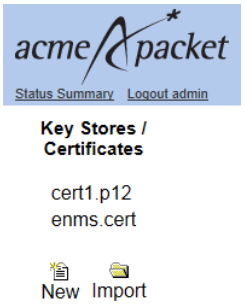

2. Enter a name and passphrase if creating a new key store and click **Create**.

   The key store appears.

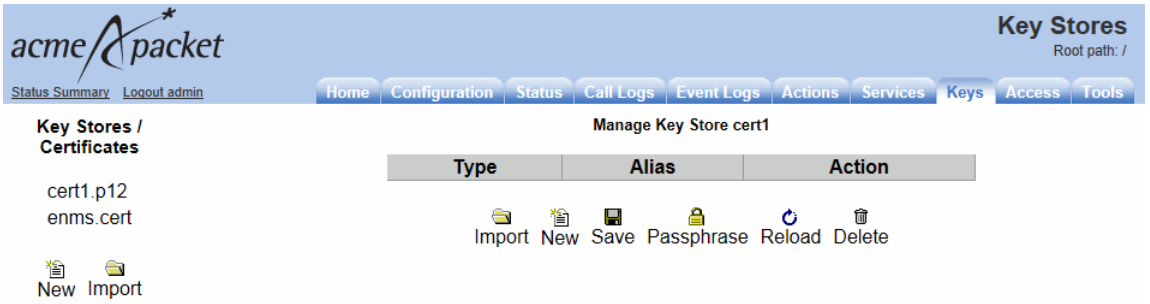

3. Click **New**. The Generate New Self-Signed Certificate in Key Store *X* page appears.

4. Click **Add** beside the **Alternate name** field to add alternate host names be added to the certificate's subjectAltName field.

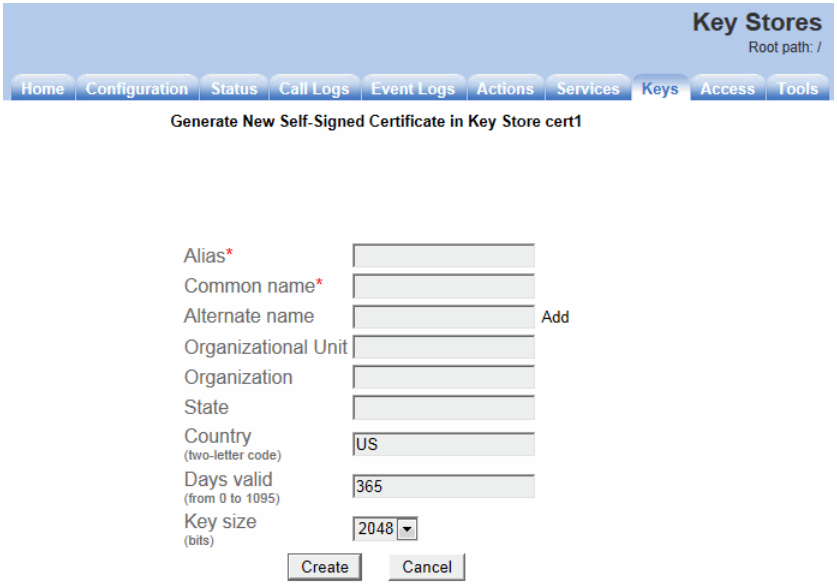

5. Click **Create**. The certificate appears in the key store.

   Note: When configuring the OS-E via the CLI, separate multiple SAN entries using the '|' character.

To view the SANs within a certificate click **View** next to the certificate name. The following image shows three SANs.

Properties	
Property	Value
Type	X.509
Version	3
Subject	CN=sjones@acmepacket.com, C=US
Issuer	CN=sjones@acmepacket.com, C=US
Valid After	Apr 29, 2013 8:59:26 AM
Valid Until	Apr 29, 2014 9:00:16 AM
Serial Number	1367240416880
Signature Algorithm	SHA1WithRSAEncryption

Other Properties	
Property	Value
Subject DN	C=US,CN=sjones@acmepacket.com
Subject Alternate Name	DNS Name:sjones@ap.com
Subject Alternate Name	DNS Name:sj@ap.com
Subject Alternate Name	DNS Name:sjones@acmepacket.com
Issuer DN	C=US,CN=sjones@acmepacket.com

Configuring the Certificate on the OS-E

Once you have imported the certificate to a directory on the OS-E system, configure the settings that control how the OS-E uses the certificate.

CLI session

The following CLI session sets the directory and certificate destination file name path and specifies the passphrase.

```
NNOS-E> config vsp
config vsp> config tls
config tls> config certificate myNetworkCert.pfx
Creating 'certificate myNetworkCert.pfx'
config certificate myNetworkCert.pfx> set allow-null-cipher enabled
config certificate myNetworkCert.pfx> set passphrase-tag pass
```

By default, the OS-E only supports TLSv1 and higher. Specify the **passphrase-tag** associated with the certificate file. Use this property if the certificate file is encrypted to have its private key information protected. This **passphrase-tag** must match the string with which the certificate was encrypted.

Displaying the Certificates Installed on the OS-E

Use the **show certificates** to command to display the list of installed certificates on the system.

Other TLS Certificate Settings

Using Certificate vs. Default-Outgoing-Settings

The OS-E uses a certificate configuration to identify the certificate file and the characteristics of the certificate. There are two types of certificate configuration—a named certificate entry that can be applied to specific TLS connects and a default certificate settings for use when a specific entry was not identified.

The entry created by the **certificate** object is used when the OS-E functions as a server in a TLS connection. Or, it can be used in an OS-E-as-client setup, if you have configured the connection to use a specific certificate. For example, when you set the connection type to the LDAP server to TLS in the **directory** object, you are required to enter a named certificate.

The entry created by the **default-outgoing-settings** object is used when the OS-E is a client with an unspecified certificate. For example, if you set the protocol that the DNS resolver server uses to TLS, you are not prompted for a certificate name. In this case, the OS-E uses either:

- The certificate identified in the **sip-settings** object, if the session matched a configured policy.
- The **default-outgoing-settings** if the session did not match a configured policy or the policy did not have a certificate specified.

Refer to the *Net-Net OS-E – Objects and Properties Reference* for detailed information on the default-outgoing-settings object under TLS.

Verifying Peer Certificates

The OS-E allows you to verify peer certificates. By default, the OS-E accepts all peer certificates. However, you can configure the OS-E to reject a connection if a peer's certificate does not meet the requirements of the network. Basic verification checks that the certificate's chain is valid, that it was signed by a trusted CA, and that the certificate has not expired.

To verify a peer's certificate, the appropriate CA file must be installed on the OS-E. For example, to connect to an LCS server, there are four requirements,

1. A client certificate that Session Presents presents to LCS at connection time,
2. A CA file (in PEM or PKCS#12) to verify the server's certificate when it is presented to the OS-E,
3. A Certificate Revocation List (CRL) in PEM format, a list of certificates that a CA has revoked, and thus can no longer be trusted. If any of the certificates in the chain presented to the OS-E appear in the CRL, the OS-E rejects the connection. This is an optional step. And,

4. A valid, verifiable host name that is presenting the certificate. If the host name doesn't match what the OS-E expects, the OS-E rejects the connection, even if the chain is valid.

CLI Session

The following CLI session defines multiple default CA files, and multiple default CRL files:

```
NNOS-E> config vsp tls
config tls> config default-ca
config default-ca> set ca-file /cxc/certs/ca1.pem tag1
config default-ca> set ca-file /cxc/certs/ca2.pem tag2
config default-ca> return
config tls> config default-crl
config default-crl> set crl-file /cxc/certs/crl1.pem tag3
config default-crl> set crl-file /cxc/certs/crl2.pem tag4
config default-crl>
```

Enabling Peer Certificate Verification

The **peer-certificate-verification** property allows you to control whether the OS-E validates a peer's certificate.

```
NNOS-E> config vsp tls
config tls> config certificate myNetworkCert.pfx
config myNetworkCert.pfx> set peer-certificate verification {none |
   if-presented | required}
```

None—The OS-E will not request a certificate from the peer, and will verify a certificate if presented with one. This is the default setting.

IfPresented—If a peer presents a certificate, the OS-E verifies it, or rejects the connection if the certificate fails verification. If no certificate is presented, the OS-E allows the connection.

Required—If a peer presents a certificate, the OS-E verifies it, or rejects the connection if the certificate fails verification. If no certificate is presented, the OS-E rejects the connection.

Note: TLS treats clients (initiators) and servers (answerers) differently. In a typical TLS connection, only the server presents a certificate; the client is only allowed to present a certificate if it is requested to do so by the server. Therefore, the **IfPresented** option applies only for a client connection.

Controlling the CA Files and CRLs to Apply to the Certificate

Configure each certificate entry to use or ignore the default CA and CRL settings.

```
config> config vsp tls certificate myNetworkCert.pfx
config certificate myNetworkCert.pfx> set use-default-ca true
config certificate myNetworkCert.pfx> set use-default-crl false
```

Optionally, you can configure each certificate entry to use an extra CA and an extra CRL, independent of the default settings, using the file path and passphrase tag.

```
config certificate myNetworkCert> set specific-ca-file /cxc/certs/
   ca9.pem tag9
config certificate myNetworkCert> set specific-crl-file /cxc/certs/
   crl9.pem tag10
```

Setting the Required Peer Name

The **required-peer-name** property specifies the name that appears in the presented certificate.

- If you do not configure a peer name, then the OS-E does not check the presented name.
- If you do configure a peer name, then that name must appear in the DNS field of the **alternateName** field, or in the **commonName** field for the certificate.

The **required-peer-name** can include wildcards, such as "*.acmepacket.com". If the presented name does not match the required name, the OS-E rejects the connection.

Configuring TLS on Ethernet Interfaces

Referring to the network illustrated in the “Steps to Configuring TLS,” section, note that one Ethernet interface is connected to public Internet on port 443, and other Ethernet interface connects to the enterprise or service provider’s private network on the known TLS port 5061. Using port 443 on the public side of the network allows HTTPS requests to pass through the network firewall to the OS-E system.

CLI Session

The following CLI session configures IP on the public and private OS-E interfaces, and the SIP protocol, ports, and TLS certificate destination name references.

```
NNOS-E> config cluster
config cluster> config box 1
config box 1> config interface eth0
config interface eth0> config ip private
Creating ‘ip private’
config ip private> set ip-address static 10.1.1.1/24
config ip private> config sip
config sip> set udp-port 5060
config sip> set tcp-port 5060
config sip> set tls-port 5061
config sip> set certificate vsp tls certificate myNetworkCert.pfx
Creating ‘vsp\tls\certificate myNetworkCert.pfx’
config sip> return
config ip public> return
config interface eth0> return
config box 1>

config box 1> config interface eth1
config interface eth0> config ip public
Creating ‘ip public’
config ip private> set ip-address static 216.1.1.1/24
config ip private> config sip
config sip> set udp-port 5060
config sip> set tcp-port 5060
config sip> set tls-port 443
config sip> set certificate vsp tls certificate myNetworkCert.pfx
Creating ‘vsp\tls\certificate myNetworkCert.pfx’
config sip> return
config ip public> return
config interface eth0> return
config box 1>
```

Chapter 6. Configuring Secure Media (SRTP) Sessions

About This Chapter

This chapter provides information on configuring inbound and outbound encryption on SIP media sessions anchored by the OS-E.

Anchoring Media Sessions

Media anchoring forces the SIP media session to traverse the OS-E system. The **auto** setting enables conditional anchoring where the OS-E uses its auto-anchoring algorithms to determine anchoring necessity based on a variety of criteria, including whether you have configured smart anchoring via the **autonomous-ip** object and whether the calling devices are behind a firewall.

The following image shows an OS-E Management System session where you enable media anchoring in the default-session-config.

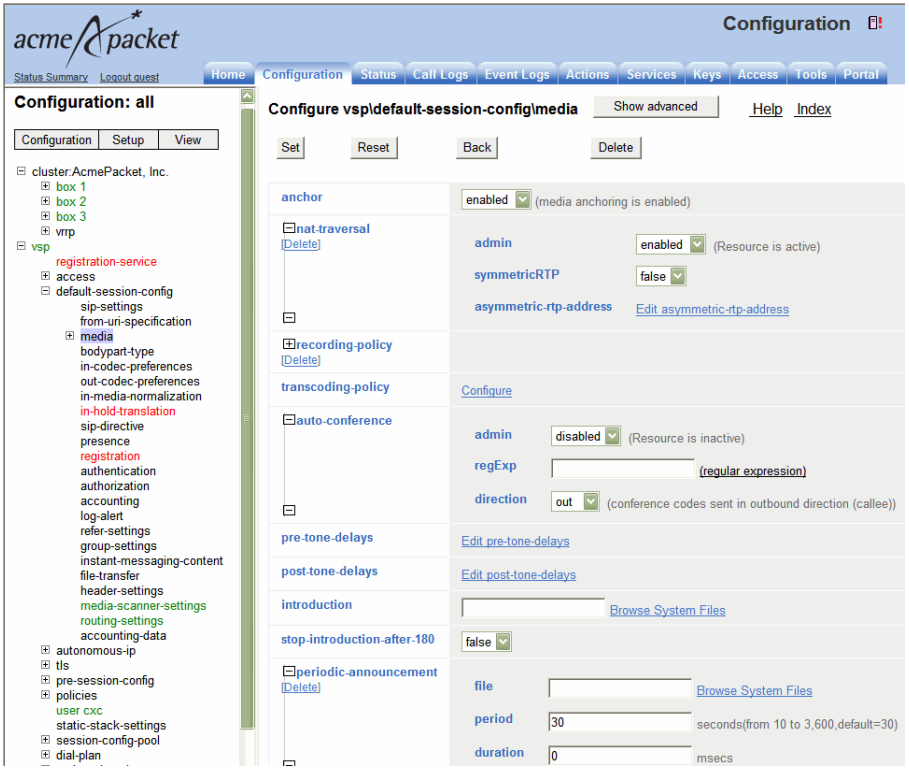

Configuring Inbound and Outbound Encryption

For secure inbound and outbound media sessions, you need to configure OS-E **in-encryption** and **out-encryption** settings. Inbound encryption handles the portion of the call from the initiator to the OS-E using a specified encryption method. Similarly, outbound encryption handles the portion of the call from the OS-E to the call recipient using a specified encryption method.

The following image shows the inbound encryption configuration page.

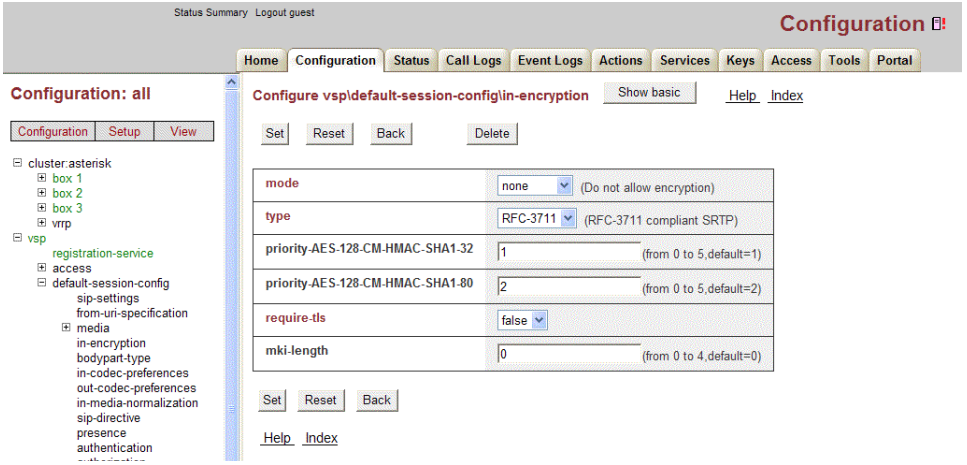

Inbound Encryption Mode and Type

Set the inbound encryption mode to one of the following settings:

- **none**—The OS-E disables the encryption put forth by the incoming endpoint. (That is, it responds “no” to the encryption portion of the authentication handshake.) If the outbound endpoint requires encryption, then the call is dropped.
- **allow**—The OS-E passes the call through, leaving the encryption setting unchanged.
- **require**—The call must come in with encryption specified or the OS-E drops it.

Set the inbound encryption type to one of the following settings:

- **RFC-1889**—Use encryption as defined in RFC 1889, RTP: A Transport Protocol for Real-Time Applications. This mode is used for compatibility with Windows Messenger and Microsoft Office Communicator, neither of which currently support RFC-3711 encryption. Instead, it uses a DES-CBC encryption of the entire UDP payload (including RTP headers) with no authentication.
- **RFC-3711**—Use encryption as defined in RFC 3711, The Secure Real-time Transport Protocol (SRTP). This is the same encryption as used in the OS-E setting.

- **Linksys**—Use Linksys/Sipura encryption over Linksys phones. Refer to Lynksys Encryption for more information

The following image shows the inbound encryption configuration page.

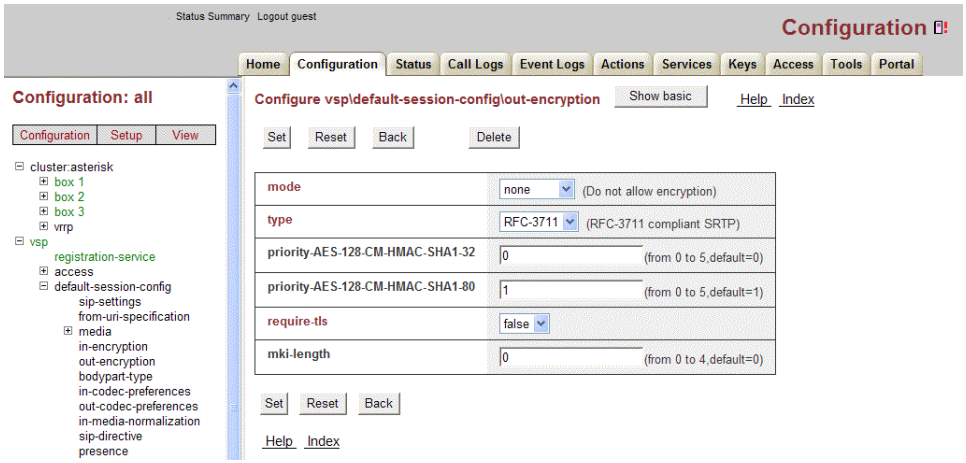

Outbound Encryption Mode, Type, and Require-TLS Setting

Set the out-encryption mode to one of the following settings:

- **none**—The OS-E disables the encryption put forth by the outbound endpoint. (That is, it responds “no” to the encryption portion of the authentication handshake.) If the inbound endpoint requires encryption, then the call is dropped.
- **offer**—The OS-E changes or establishes the encryption type to the value specified in the **type** property, below.
- **follow**—If the inbound endpoint offered encryption, the OS-E offers that type to the outbound endpoint.
- **require**—The call must come in with encryption specified or the OS-E drops it.

Set the out-encryption type to one of the following settings:

- **RFC-1889**—Use encryption as defined in RFC 1889, RTP: A Transport Protocol for Real-Time Applications. This mode is used for compatibility with Windows Messenger and Microsoft Office Communicator, neither of which currently support RFC-3711 encryption. Instead, it uses a DES-CBC encryption of the entire UDP payload (including RTP headers) with no authentication.

- **RFC-3711**—Use encryption as defined in RFC 3711, The Secure Real-time Transport Protocol (SRTP). This is the same encryption as used in the OS-E setting.
- **Linksys**—Use Linksys/Sipura encryption over Linksys phones. Refer to Linksys Encryption for more information.

Note: Because the OS-E does not always know on the outbound leg the encryption method expected by the recipient (because that recipient isn't in the registry), you must manually set the type of encryption to offer.

Require TLS

The **require-tls** property specifies the requirements of the signaling protocol for a call's outbound leg. That is, it defines whether the OS-E offers SRTP over a non-secure (TCP or UDP) signaling connection. The action of this property depends on the setting of the mode property. When this property is set to:

- **true**—The OS-E only offers encryption when talking to a TLS client. If TLS and SRTP are required (**mode** is set to **require**), the OS-E fails calls going to TCP/UDP clients. If the mode property is set to **offer** or **follow**, the OS-E forwards the call without SRTP.
- **false**—The OS-E offers SDP messages according to the mode setting without regard for the signaling transport. This allows keys to be exchanged in an insecure message.

Most phones follow RFC 4568, SDP Security Descriptions for media Streams, and thus require that this property be set to *true*.

Linksys Encryption

The **linksys** action allows you to generate a Linksys/Sipura mini-certificate and private key which, after loaded into the phone, will be used to exchange the symmetric key. You must execute this action and load the result into both phone parties.

Linksys equipment supports a proprietary version of SRTP. It uses SIP INFO messages to exchange credentials (in mini-certificates) and securely distribute the key used to encrypt/decrypt the RTP packets. The RTP encryption is a variation of RFC-3711; the encryption algorithm is the same (AES-CM-128), but uses HMAC-MD5 instead of HMAC-SHA1 for authentication.

The CLI syntax for the **linksys mini-certificate** action is:

```
linksys mini-certificate user-id display-name expires [filename]
```

The following image shows the linksys mini-certificate page.

The **linksys** action provides three tools:

- **mini-certificate**—Creates a mini-certificate, which will later be used by a Linksys phone to exchange an encrypted symmetric key. When both phones in a call support cryptographic exchange, use this action to create a mini-certificate that is sent in an INFO message to the other phone. (You must execute this action for both phones.) After exchanging mini-certificate, the phones can then exchange an encrypted symmetric key.

  Enter the following fields to generate a mini-certificate:

  - **userID**—A name that identifies this phone (subscriber) to the other party. The user ID can be up to 32 characters.
  - **displayName**—A name used by the caller to verify that the callee is the intended call recipient. Enter the user ID field in the Request URI of the INVITE message sent to the proxy server by the caller UAC when making a call to this subscriber (UAS). The display name can be up to 16 characters.
  - **expiration**—The date and time at which this mini-certificate expires. Enter the date in the format *hh*:*mm*:*ss yyyy*-*mm*-*dd*.

- **filename**—A name for an output file that will contain the mini-certificate and private key. If you do not specify a file name, the output is not written to a file.

Once you execute this option, the OS-E returns the content of the mini-certificate and the SRTP private key. You can copy and paste each of these fields into your phones Web GUI (or other software interface), as well as test the certificate using the **check-mini-cert** option.

- **generate-ca-key**—Generates a Linksys/Sipura CA key. This is the public/private key pair that acts as the Sipura certificate authority. It is needed to generate the mini-certificates for each phone and during the key exchange.

  The key is stored in **/cxc/certs/linksys_ca.pem**. When executing this action, you can specify whether to overwrite any previous CA key. The default setting, **false**, does not overwrite the key. Set the field to **true** to force an overwrite.

- **check-mini-cert**—Verifies the contents of a certificate created with the mini-certificate option. When executed, the OS-E checks the expiration date and signature of the certificate. Enter the content of the mini-certificate to invoke this option.

Note: You must have a root certificate loaded on the OS-E system for this action to be successful. The default location for the root certificate is **/cxc/certs/linksys_ca.pem**.

The following CLI session generates the mini-certificate and private key for the Linksys phones.

CLI Session

```
NNOS-E> linksys mini-certificate ?

Generate or load Linksys/Sipura SRTP encryption parameters

 syntax: linksys mini-certificate user-id display-name expires
   [filename]
         linksys generate-ca-key [force]
         linksys check-mini-cert minicert

NNOS-E> linksys mini-certificate 9577 9577-display "23:05:45
 2004-11-25"

Mini Certificate:
```

```
OTU3Ny1kdWVyb2QAAAAAAAAAAAAAAAAAAAAAAAAAAA5NTc3AAAAAAAAAAAAAAAAMDMwNT
 Q1MjYxMTA0ybYgcwG8IeaYz225Grs7sDJflnfyJxARPehQ+CO6WisAZ77U2zBi8TCapI
 wqcDhNXwgYKZxljAET3dFnzAxs2ze1/
 kEHCqvUmDIEjaYL+1WTySaI1TGKy15FbyZb6dQXtbPF+fXiRP//
 caFfKUBTuuwtjExxaAz0H3u8Tc2YT/wH7a0+snpUTFeK/
 Sv9vd7aAUbufSxewlL2GeTdOu0v2i4R25/
 RH6iOHyChGpVt2EJ3BHAlLgXTfJibiwwkrMSe1grSibsCy0D825ezAt66AVKTa/
 hOmSBvdZvdamJIsbP89vnAJPiOfWNet8T40/wOYyylAE5JDJ/2+G/
 MDyc5ImzFTvifKvIQ55T7Jr5E0RUbacDZIlHy5oW+x4sfawCiQZunnb11qlAgYhvOeuo
 4f3JGUKJAld0GRjHfvjRhb3c=

SRTP Private Key:
Oxq38oJqjhe++yBTtTotoMndnZXulkgnnxFQPd0v96oc81IZ5dug9Szob9ZYQXsPkWAxSb
Oxq38oJqjhe++BVpyxz2P2qtZEg==
NNOS-E> linksys check-mini-cert
 OTU3NwAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA5NTc3AAAAAAAAAAAAA
 AAAMTU1ODQ4MTEyNTA0z1TBkpXzjmR6PFX5K4S7G5SxdpozH460T14KpwOxZ8
 ly4KWpFlcC2rTTWEU6WnOufcj5Bfif7cdsAF/
 89kZu83NFceK2ZBRGrJ4cbxREtuPwy1FqkXpBQcztTFXjeyFaq8K7OESebQay
 FetBEceIupuzxfedlJPRsMRhsHN1uKpomc/
 tdJFHJhxSzn+fX+GTACrXQEHzI+ooDL+iQvzhJ1zk/
 gXTGuk76lkJG2XLvSvdjTp8RjQX/F5h0GnBa02d3bQ51n7IBvJnTeaGKp/U/
 e5pQvW5u6vD/uHkqkTGkZDZzOyIISIdgWVxdjA9cpaSa2D5nPhr8G/
 WhOadLZ08fmB0kPwEFjJ0h0dojjknjNJp/
 qVjR5NEEzuj5kH7Qlvxk25l0MThhydCYpbxShy2GSno7apnyCA02YBQCRlGBO
 s=
Certificate has expired

NNOS-E> linksys generate-ca-key
Unable to overwrite Linksys CA key
```

Chapter 7. Creating and Commissioning USB Sticks

This chapter provides information on creating and commissioning OS-E USB software installation sticks for commissioning third-party servers. Using the Internet and secure Web URLs, Oracle provides all necessary software downloads for USB creation, product licensing, and commissioning of the OS-E on your selected hardware.

As part of each download, and depending on your actual requirements, Oracle can provide the following:

- USB Boot Media Creator (BMC) with the OS-E software.
- Documentation on how to create a USB stick and commission the OS-E software on your selected hardware
- Standard set of Oracle OS-E technical publications

Supported USB Sticks

You must provide a USB stick with at least 1GB storage, to handle OS-E software downloads. Oracle has tested a variety of USB sticks available from current suppliers and manufacturers. Most 4GB USB sticks manufactured today work.

USB Stick Restrictions

Files that are larger than 2 MB will not be backed up to the USB stick and restored during the upgrade process.

All *.cfg and *.xml files in the current working directory (/cxc) that are less than 2 MB in size are backed up to the stick and restored during the upgrade.

In the event that Internet access is unavailable, use the **show system-info** command to display the box identifier. Access https://licensecodes.oracle.com and click Acme Packet. The Acme Packet License Key Request page appears. Enter the required fields, including the **show system-info** value for **Original Chassis Serial #(s)**, and click **Submit**. After your customer information has been verified, Oracle will send the license(s) to you via email. You can then use the OS-E Management System **Upload License** function or **WinSCP** to place the license files on the system.

Note: Licenses are applicable to royalty-bearing codecs only.

- **Downloads for the OS-E Virtual Machine**—Includes the VM for running the OS-E/ASC software on x86-based hardware running your VM. The OS-E and ASC technical documentation download is also included.

  For complete information on the VM, see "Installing and Running the OS-E Virtual Machine".

Important Note About the New USB Stick

The OS-E USB stick provides three important functions:

- **Rescue utilities**—After you have successfully commissioned the system (booted and licensed), the system automatically rewrites the USB stick so that you can use it to run system utilities in the event of a catastrophic failure. You will not be able to use this USB again to license another system. The USB can only be used at the specific system from which it was originally written.

  Additionally, licensing information is rewritten to the USB, directly associating the license with the system. Use the **show system-info** action to display the box identifier (box-id) to which this USB is associated. The USB also contains log and debug files that you can use to help diagnose problems associated with the USB licensing process.

- **Rescue stick creation**—With the original USB commissioning stick, use the **restore-stick-create full-backup** action to capture the current software, certificates, and operating system image to the USB stick. Oracle recommends that you use **restore-stick-create full-backup** to preserve the image prior to performing a system software upgrade, or whenever you have made significant and reliable changes to the system configuration.

See Chapter 3 of this guide for information on these functions.

Creating a New USB Rescue Stick

The **restore-stick-create full-backup** action allows captures the current OS-E software and the operating system image and creates a new USB rescue stick. Oracle recommends that you use **restore-stick-create full-backup** to preserve the image prior to performing a system software upgrade, or whenever you have made significant and reliable changes to the system configuration.

Perform the following steps:

1. At the NNOS-E prompt, type **umount usb**.

   Ignore the warning about the USB stick not being mounted.

2. Remove the USB stick and wait at least five seconds before reinserting the stick.
3. Invoke the **restore-stick-create full-backup** action. The resulting USB is a boot device from which you can restart and restore the OS-E.

Note: The log event indicating that the operation has completed successfully appears while data is being written to the stick. DO NOT immediately remove the USB stick when you see this log event. Instead, issue the **umount usb** command again, and wait for it to complete.

4. Remove the restore USB when the **restore-stick-create** action has completed.

Note: PostgreSQL database, media recordings, system tar files(.gz) are not written to the restore stick with the **restore-stick-create** action.

Note: Use the **restore-stick-create config-backup** action to create a restore stick containing the current configuration file only.

Using the Rescue Utility USB

To use the rescue utility USB, perform the following steps:

1. Insert the rescue utility USB into one of the USB slots.
2. At the NNOS-E prompt, if available, enter **restart cold** to do a full restart.

3. The USB and disk lamps will blink during the boot-up process followed by a series of system messages. Press any key to continue, or perform the appropriate action below. On the

   - **Sun Netra X5-2**—Press **F8**
   - **Sun Netra X3-2**—Press **F8.**
   - **HPDL360 G7**—Press **F11**.
   - **HPDL585 G7**—Press **F11**.
   - **HPDL320 G8**—Press **F11**.
   - **HPDL360 G8**—Press **F11**.
   - **HPDL160 G9**—Press **F11**
   - **CXC-350** — Press the **ESC** key.
   - **Sun x4150** — Press **F8**.
   - **IBM X3650, X3550, X-3350** — Press **F12**.
   - **Dell 1950 and 2950** — Press **F11**.
   - **Dell R220** — Press **F11**.
   - **IBM HS20 and HS21** — Press **F12**.
   - **Fugitsu Siemens RX200 S5** — Press **F12**.
   - **Fugitsu Siemens RX300 S4** — Press **F12**.
   - **ATCA MCLB0040** — Press **F8**.

4. The system enters utility mode and the following menu appears:

You can perform one of the following:

- Rescue an Oracle Linux system

  Note: This boots the system into rescue mode and attempts to mount the OS-E drive under /mnt/sysimage.). Refer to Using the Rescue Mode for more information.

- Install Oracle Linux 7.2.
- Test this media and install Oracle Linux 7.2.
- Install Oracle Linux 7.2 in basic graphics mode.

  Note: The above three options run the Oracle Linux 7.2 installer. These actions wipe out any existing configuration and data. Using either the **restore-stick-create** action or the rescue mode, ensure your configuration and other data files have been backed up. Once Oracle Linux is reinstalled, the OS-E packages must be reinstalled. For more information on installing the OS-E, see "Installing the OS-E".

- Run a memory test.
- Boot from the local drive. (Boots from the local hard disk and defaults to the first disk. To change the disk, use the **<*Tab*>** key.

Using the Rescue Mode

The Expert mode provides utilities that allow you to recover from system failures where there is apparent damage to the software and configuration file, and where recovery is necessary to return the OS-E to normal operation.

If you enter rescue mode, the system boots and attempts to mount the OS-E hard drive at /mnt/sysimage. If there are multiple Oracle Linux installations, there is a prompt to select which installation.

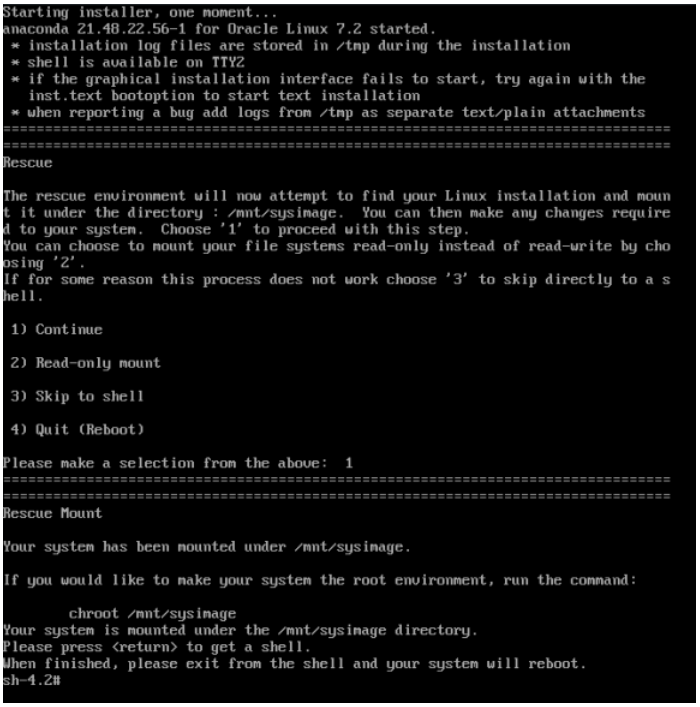

If the system mounts the drive successfully, you can attempt to repair it or perform a backup.

Backing Up the Configuration

Execute the *save_box_config.sh* script to copy the configuration.

Note: The USB drive must be mounted prior to executing this script.

```
sh-4.2# mkdir /mnt/usb
sh-4.2# mount /dev/disk/by-label/ASC_RESCUE /mnt/usb
```

```
sh-4.2# save_box_config.sh ?mnt/sysimage /mnt/usb
sh-4.2# umount /mnt/usb
```

After a fresh Oracle Linux and OS-E installation, booting with this USB stick inserted automatically restores this configuration.

System and Data Drive Locations

The following system and data drive locations are for Oracle Net-Net OS-E hardware. For all other supported platforms and third-party servers, refer to the documentation that accompanies the hardware.

- **NN 2610**— System drive is disk 1 (left slot); remaining drives (2) are data drives.
- **NN 2620**— System RAID-1 drives are disk 1 (lower left slot) and disk 2 (upper left slot); remaining RAID-10 drives (4) are data drives
- **CXC-50** — Single disk only for both system and data.
- **CXC-350 and CXC-354** — System drive is disk 1 (left slot); remaining drives (2) are data drives.
- **CXC-550** — System drive is disk 1 (top left slot); data drive is disk 2 (lower left slot).
- **CXC-554** — System RAID-1 drives are disk 1 (lower left slot) and disk 2 (upper left slot); remaining RAID-10 drives (4) are data drives.
- **CXC-1250** — System drive is disk 1 (left slot); disk 2 is a data drive.

Installing and Running the OS-E Virtual Machine

This chapter provides information on downloading, installing, and running the OS-E Virtual Machine (VM) software in virtual OS environments. This software is the same software as used for non-virtual OS but has been packaged specifically as a VM for use in virtual OS environments.

The OS-E VM is designed to be used as an evaluation platform so that potential customers can test the OS-E software in an environment that does not require them to install the software on a dedicated piece of hardware. In some cases, the VM can also be used in production environments provided that the customer understands the limitations associated with using the VM software in a virtual OS environment.

Server-Based Requirements

Before downloading the VM to an x86-based server, make sure that you have met the following hardware and software requirements:

Hardware

- x86-based Windows or Linux server with Intel 32- or 64-bit dual-core processors
- 2GB minimum (4 GB recommended) physical memory for each VM instance
- Minimum of 40GB hard disk space per VM instance
- One or two Ethernet interfaces

Software

The following VM platforms have been certified for use with the OS-E:

- OVM 3.3.1
- VMware ESXi 5.5
- XEN 3.4.3

Linux Installations

If you are installing the OS-E VM on a Linux workstation running VMware, Oracle recommends the following technical resources:

For Server 1.0

http://www.vmware.com/support/pubs/server_pubs.html

Player 1.0 and 2.0

http://www.vmware.com/support/pubs/player_pubs.html

Installing the VM

This section describes the process for installing the OS-E VM on each of the certified VM platforms.

Installing the OS-E on an Oracle Virtual Machine

The OS-E is certified to run on the Oracle Virtual Machine (OVM) 3.3.1.

Prerequisites

You must meet the following prerequisites before installing the OS-E on an OVM:

- A Network File System (NFS) has been mounted for VM storage with an additional storage file server for repository
- A server pool has been created
- Server(s) have been discovered and added to this pool
- The ISO file has been imported

- Networks and Virtual MAC range have been created
- VM Console access (VNC) has been made available

You create the OS-E VM via the OVM Manager GUI. The OVM Manager binds to the weblogic server on the Oracle Linux host's 7002 SSL port.

Access the OVM Manager using the following link:

`https://x.x.x.x:7002/ovm/console`

Where *x.x.x.x* is the OVM Manager's IP address.

1. Log into the OVM Manager using the user name and password configured when you set up the OVM.
2. Create external routable interfaces by selecting the **Networking** tab, selecting the **Networks** button, and clicking the plus (+) icon.

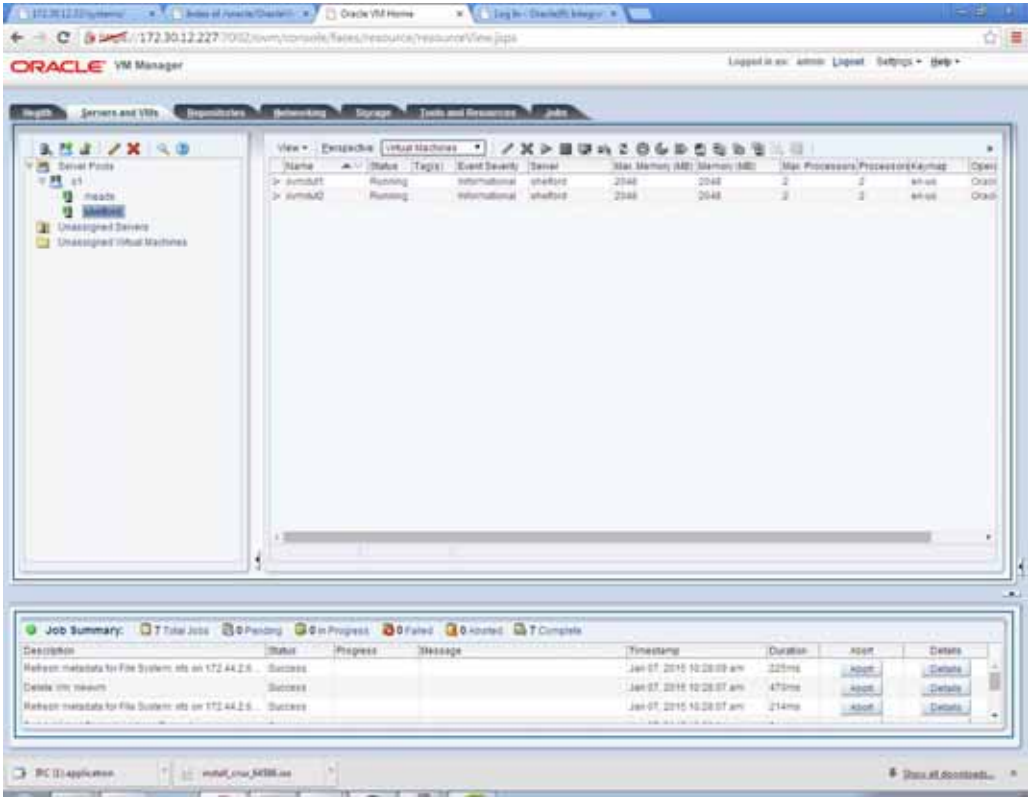

3. Create a new bridge **bonds/ports only** and select **Virtual Machine** in the **Network Uses** field.
4. Bind the new bridge to a free port on the VM host.
5. *Optional*. Create a heartbeat interface (if you choose to configure clustered VMs) by selecting the **Networking** tab, selecting the **Networks** button, and clicking the plus (+) icon.
6. Create a new bridge local network only and select **Virtual Machine** in the **Network Uses** field.
7. Create MAC addresses for each Virtual NIC by selecting the **Networking** tab, selecting the **Virtual NICs** button, and creating a **Dynamic MAC Address Range**.

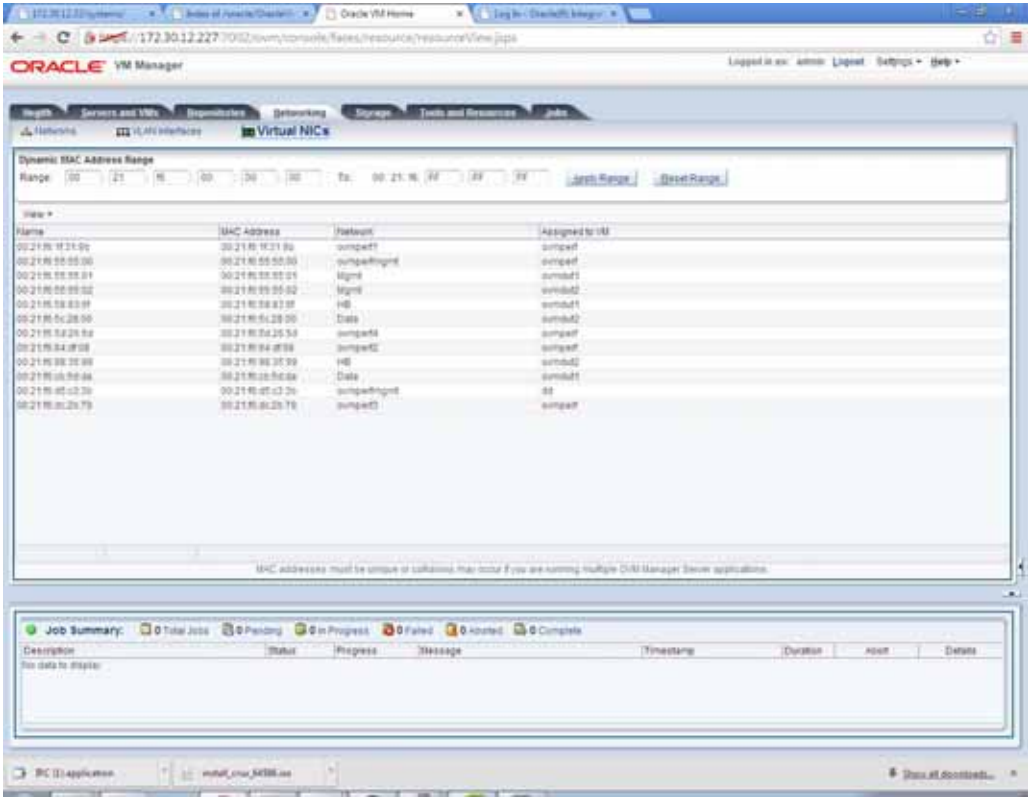

Note: You must create a unique MAC address for each Virtual NIC.

As mentioned previously in Prerequisites, you must mount an NFS to host the OS-E. The following image shows a properly configured NFS mount point for the VM.

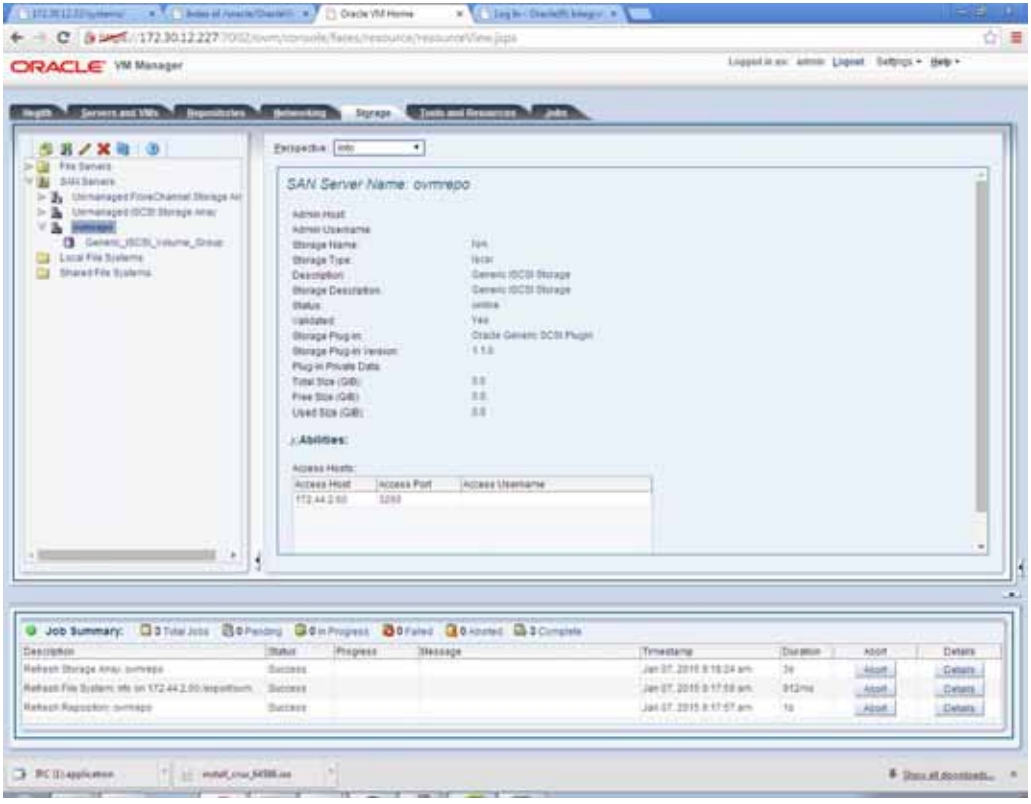

When creating a VM, your storage repository contains an ISO file and the new VM immediately boots from the Virtual DVD.

To boot the VM from the Virtual DVD.

1. Select the **Repositories** tab and select the repository you created from the **File Server**.
2. Select **ISOs**.

3. Via HTTP, import the OS-E's .iso code.

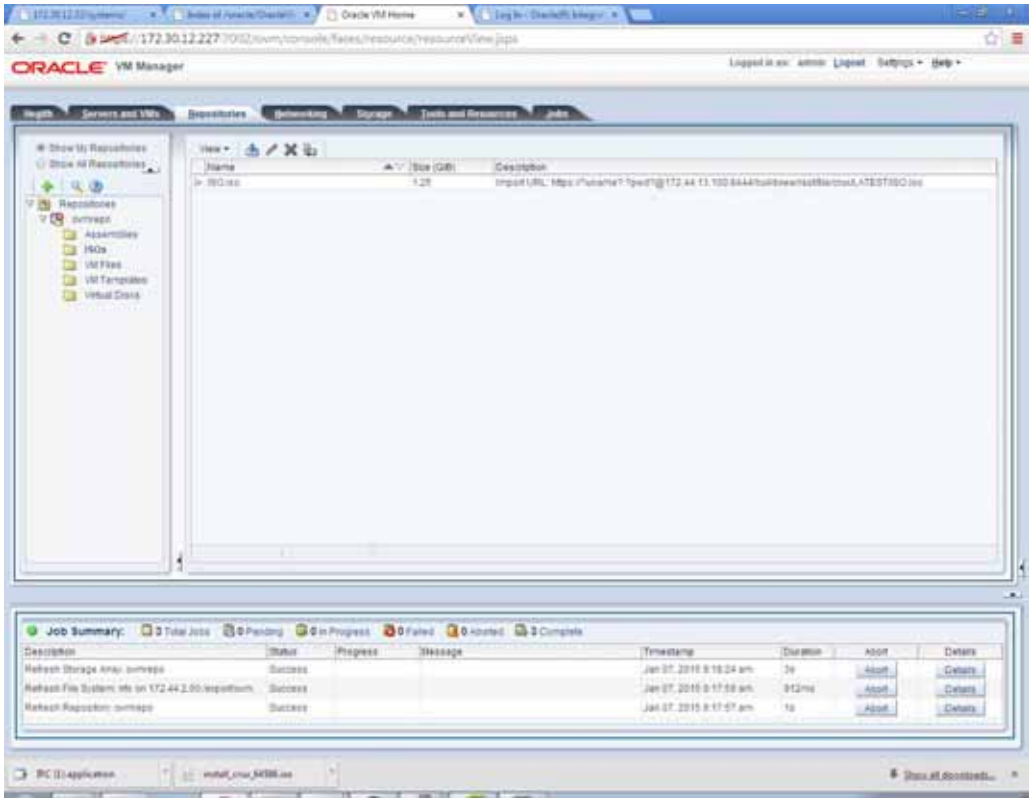

You are now ready to create a VM.

To create a VM:

1. Select the **Servers and VMs** tab and choose the server on which you are hosting the VM.

2. Select **Create Virtual Machine**.

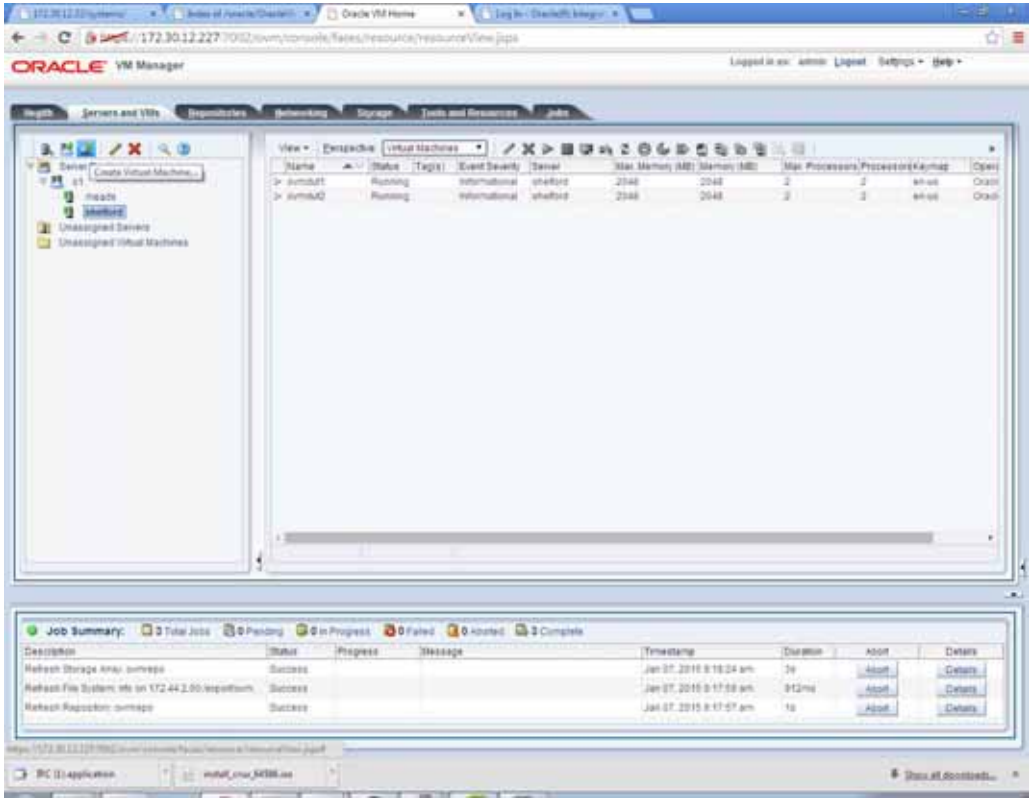

3. Click **Next**.

4. Specify a **Name** and set the **Memory** and **Processors** for this VM.

Note: The default **Memory** is **1024** and the default number of **Processors** is **1**.

5. Click **Next**.
6. Select your networks.

   Note: The order you select the networks affects how the OS-E ethernets align.

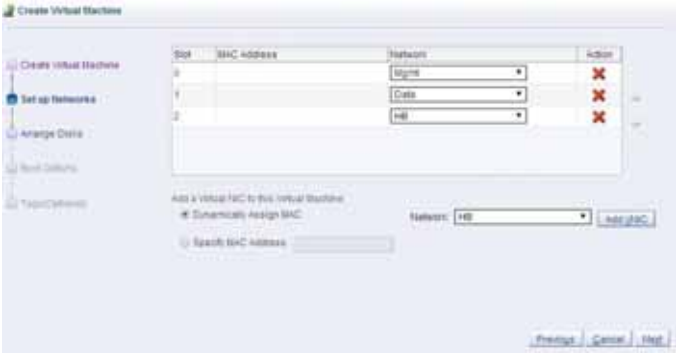

7. Click **Next**.

   These MAC addresses (whether assigned dynamically or statically) now appear as assigned MAC addresses under the Virtual NICs tab in OVM Manager.

To create the VM virtual disk:

1. Select the Virtual Disk's **Disk Type** and select the **Create a Virtual Disk** icon.

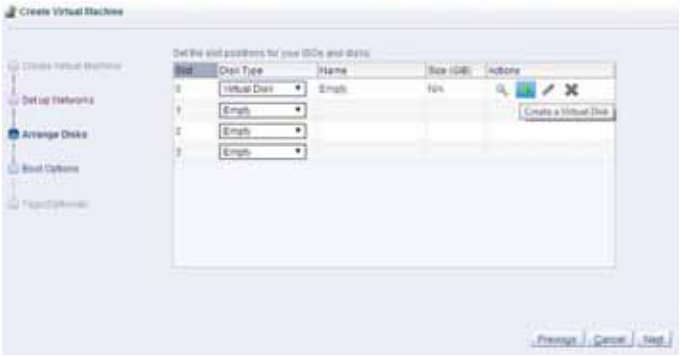

2. Click **Next**.
3. Select the previously-created **Repository** and enter a **Virtual Disk Name** and a **Size** (Oracle recommends 40 GB).

4. Click **OK**.

To point the OS-E ISO code to commission the VM:

1. Select **CD/DVD** from the Slot 2 **Disk Type** drop-down menu and select **Select an ISO**.

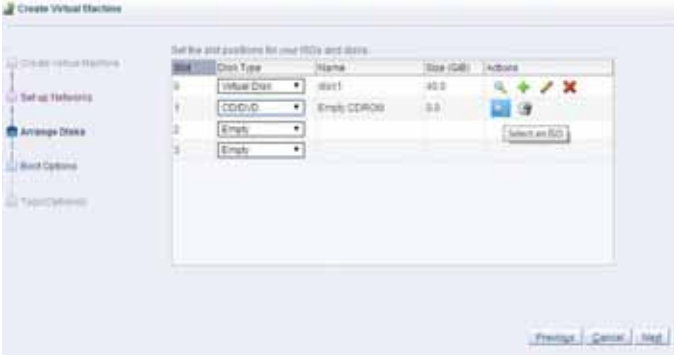

2. Click **Next**.
3. Select the previously-imported ISO..

4. Click **OK**

Once the OS-E ISO code is pointed to commission the VM, set the Boot Options. The first time you boot you utilize the CDRom as nothing resides on the Disk yet. All subsequent boots utilize the Disk and ignore the CDRom.

> Note: If you choose CDRom as the first boot option, the initial boot, as well as all subsequent boots, continue to utilize the CDRom.

To set the Boot Options:

1. Select **Disk**.
2. Select **CDRom**.

3. Click **Finish**.

To see the newly-created VM, select the **Servers and VMs** tab and click **Virtual Machines** from the **Perspective** drop-down menu. At this point in the installation process, the **Status** of the VM is **Stopped**.

To start the VM:

1. Select **Start** to start the VM.

2. Select **Launch Console**.

The OVM Console now displays the installation process. Eventually the following screen appears.

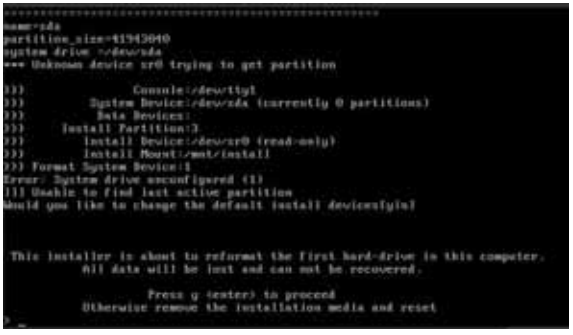

3. Type **y** and press **<*Enter*>** to complete the installation process.

The VM reboots and once the installation is complete you see the OS-E login prompt. The OS-E is now ready to be set up and configured.

Configuring OVM Passthrough

On the OVM, there are two ways to directly connect a VM to a physical port: Single Root I/O Virtualization (SR-IOV) and Peripheral Component Interconnect (PCI) Passthrough. You configure hardware passthrough at the OVM Server's CLI.

Note: Prior to configuring hardware passthrough, you must have a fully built VM, however, any NICs designated for hardware passthrough may not have an associated Network. For more information on specifying an NIC Network, see page 128.

SR-IOV is a specification that treats a single physical device as multiple separate Virtual Functions (VF)s.

Note: In development, SR-IOV was found to be available on 10GB ixgbe devices only.

To configure SR-IOV:

1. Access and log into the OVM Server's CLI.
2. Install the necessary packages on the OVM Server.

```
libibumad-1.3.8-2.mlnx1.5.5r2.el5.x86_64.rpm
libibmad-1.3.9-7.mlnx1.5.5r2.el5.x86_64.rpm
opensm-libs-3.3.15-6.mlnx1.5.5r2.el5.x86_64.rpm
kernel-ib-1.5.5.092-2.6.39_300.29.1.el5uek.x86_64.rpm
infiniband-diags-1.5.13.MLNX_20120708-4.mlnx1.5.5r2.el5.x86_64.rpm
ovsvf-config-1.0-6.noarch.rpm
```

3. Create a python script called **vnfs.py** to view and marry PCI addresses to interfaces.

```
#!/usr/bin/python
# Copyright (C) 2012 Steve Jordahl
#
# This program is free software; you can redistribute it and/or modify
# it under the terms of the GNU Lesser General Public License as
   published
# by the Free Software Foundation; version 2.1 only.
#
# This program is distributed in the hope that it will be useful,
# but WITHOUT ANY WARRANTY; without even the implied warranty of
# MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the
# GNU Lesser General Public License for more details.
#
# vfns: list SR-IOV virtual functions

import os

info = {}

def catFile(filename):
        readfile = open(filename)
        return readfile.read().strip()
```

```
for dev in os.listdir('/sys/class/net'):
        if dev.startswith('eth'):
                info[dev] = {}
               info[dev]['address'] = catFile('/sys/class/net/' + dev
   + '/address')

for dev in info.keys():
        devLink = os.readlink('/sys/class/net/' + dev + '/device')
        info[dev]['pci address'] = devLink[-7:]
        os.chdir('/sys/class/net/' + dev)
        for devInfo in os.listdir(devLink):
                if devInfo.startswith('virtfn'):
                        info[dev][devInfo] =
   os.readlink(os.path.join(devLink,
devInfo))[-7:]

for dev in sorted(info.keys()):
        print dev
        for detail in sorted(info[dev].keys()):
                print "     " + detail + ":  " + info[dev][detail]
```

4. Create **/etc/pciback/pciback.sh**.

```
#!/bin/sh
if [ $# -eq 0 ] ; then
echo "Require a PCI device as parameter"
exit 1
fi
for pcidev in $@ ; do
if [ -h /sys/bus/pci/devices/"$pcidev"/driver ] ; then
echo "Unbinding $pcidev from" $(basename $(readlink /sys/bus/pci/
   devices/"$pcidev"/driver))
echo -n "$pcidev" > /sys/bus/pci/devices/"$pcidev"/driver/unbind
fi
echo "Binding $pcidev to pciback"
echo -n "$pcidev" > /sys/bus/pci/drivers/pciback/new_slot
echo -n "$pcidev" > /sys/bus/pci/drivers/pciback/bind
done
```

5. Use an Input/Output Memory Management Unit (IOMMU) to allow both the VM and physical devices access to memory. The IOMMU allows the OVM to limit what memory a device is allowed and gives the device the same virtualized memory layout that the guest sees.

```
Edit /boot/grub/grub.conf to enable iommu and comment out the existing
   kernel entry ( see example )

# grub.conf generated by anaconda
#
# Note that you do not have to rerun grub after making changes to this
   file
```

```
# NOTICE:  You have a /boot partition.  This means that
#          all kernel and initrd paths are relative to /boot/, eg.
#          root (hd0,0)
#          kernel /vmlinuz-version ro root=/dev/sdb2
#          initrd /initrd-[generic-]version.img
#boot=/dev/sdb
default=0
timeout=5
splashimage=(hd0,0)/grub/splash.xpm.gz
hiddenmenu
title Oracle VM Server-ovs (xen-4.3.0 3.8.13-26.4.2.el6uek.x86_64)
root (hd0,0)
        #kernel /xen.gz console=com1,vga com1=57600,8n1
    dom0_mem=max:1776M allowsuperpage dom0_vcpus_pin dom0_max_vcpus=20
        kernel /xen.gz console=com1,vga com1=57600,8n1
    dom0_mem=max:1776M allowsuperpage
    iommu=passthrough,no-qinval,no-intremap
        module /vmlinuz-3.8.13-26.4.2.el6uek.x86_64 ro
    root=UUID=e2b44279-55a5-48b9-b910-82446b7b8c65 rd_NO_LUKS
    rd_NO_LVM LANG=en_US.UTF-8 rd_NO_MD SYSFONT=latarcyrheb-sun16
    KEYBOARDTYPE=pc KEYTABLE=us rd_NO_DM rhgb quiet
        module /initramfs-3.8.13-26.4.2.el6uek.x86_64.img
```

6. Add SR-IOV support to ovs.conf.

   Note: The following example configures support for 10 VFs on the server's 4 ixgbe interfaces (matching up to ethernets 9-12).

```
[root@meads ~]# vi /etc/modprobe.d/ovs.conf
options bnx2x disable_tpa=1
options ipv6 disable=1
# SRIOV support
options ixgbe max_vfs="10,10,10,10,0,0,0,0,0,0,0,0,0,0"
install ixgbe /sbin/modprobe pciback ; /sbin/modprobe --first-time
    --ignore-install ixgbe
```

7. Blacklist the Intel VF driver (ixgbevf) in dom0 so that the dom0 kernel does not try to use the VFs.

```
[root@Meads ~]# vi /etc/modprobe.d/blacklist.conf
#
# Listing a module here prevents the hotplug scripts from loading it.
# Usually that'd be so that some other driver will bind it instead,
# no matter which driver happens to get probed first.  Sometimes user
# mode tools can also control driver binding.
#
# Syntax:  driver name alone (without any spaces) on a line. Other
# lines are ignored.
#

# watchdog drivers
blacklist i8xx_tco
```

```
# framebuffer drivers
blacklist aty128fb
blacklist atyfb
blacklist radeonfb
blacklist i810fb
blacklist cirrusfb
blacklist intelfb
blacklist kyrofb
blacklist i2c-matroxfb
blacklist hgafb
blacklist nvidiafb
blacklist rivafb
blacklist savagefb
blacklist sstfb
blacklist neofb
blacklist tridentfb
blacklist tdfxfb
blacklist virgefb
blacklist vga16fb
# ISDN - see bugs 154799, 159068
blacklist hisax
blacklist hisax_fcpcipnp

# intel ixgbe sr-iov vf (virtual function) driver
blacklist ixgbevf
```

8. Reboot the OVM server.

9. Run the vnfs script to view addresses and VFs statistics.

```
[root@Meads ~]# ./vnfs.py
eth0
     address:  a0:36:9f:2c:39:74
     pci address:  30:00.0
eth1
     address:  a0:36:9f:2c:39:75
     pci address:  30:00.1
eth10
     address:  00:21:28:a1:e2:41
     pci address:  88:00.1
     virtfn0:  88:10.1
     virtfn1:  88:10.3
     virtfn2:  88:10.5
     virtfn3:  88:10.7
     virtfn4:  88:11.1
     virtfn5:  88:11.3
     virtfn6:  88:11.5
     virtfn7:  88:11.7
     virtfn8:  88:12.1
     virtfn9:  88:12.3
eth11
```

```
     address:  00:21:28:a1:e2:42
     pci address:  98:00.0
     virtfn0:  98:10.0
     virtfn1:  98:10.2
     virtfn2:  98:10.4
     virtfn3:  98:10.6
     virtfn4:  98:11.0
     virtfn5:  98:11.2
     virtfn6:  98:11.4
     virtfn7:  98:11.6
     virtfn8:  98:12.0
     virtfn9:  98:12.2
eth12
address:  00:21:28:a1:e2:43
     pci address:  98:00.1
     virtfn0:  98:10.1
     virtfn1:  98:10.3
     virtfn2:  98:10.5
     virtfn3:  98:10.7
     virtfn4:  98:11.1
     virtfn5:  98:11.3
     virtfn6:  98:11.5
     virtfn7:  98:11.7
     virtfn8:  98:12.1
     virtfn9:  98:12.3
eth2
     address:  a0:36:9f:2c:39:76
     pci address:  30:00.2
eth3
     address:  a0:36:9f:2c:39:77
     pci address:  30:00.3
eth4
     address:  a0:36:9f:2d:0b:a8
     pci address:  a0:00.0
eth5
     address:  a0:36:9f:2d:0b:a9
     pci address:  a0:00.1
eth6
     address:  a0:36:9f:2d:0b:aa
     pci address:  a0:00.2
eth7
     address:  a0:36:9f:2d:0b:ab
     pci address:  a0:00.3
eth8
     address:  00:21:28:a1:e2:46
     pci address:  5f:00.0
eth9
     address:  00:21:28:a1:e2:40
     pci address:  88:00.0
     virtfn0:  88:10.0
     virtfn1:  88:10.2
```

```
     virtfn2:  88:10.4
     virtfn3:  88:10.6
     virtfn4:  88:11.0
     virtfn5:  88:11.2
     virtfn6:  88:11.4
     virtfn7:  88:11.6
     virtfn8:  88:12.0
     virtfn9:  88:12.2
```

10. Run the module.

```
[root@meads ~]# modprobe xen-pciback
```

11. Assign devices to pciback in the format:

```
Domain 0:Bus:#:Device#:Function #).
```

Note: In the following example the 4 interfaces are VFs on ethernets 9-12.

```
[root@meads ~]# /etc/pciback/pciback.sh 0000:88:10.0
Unbinding 0000:88:00.0 from ixgbe
Binding 0000:88:00.0 to pciback

[root@meads ~]# /etc/pciback/pciback.sh 0000:88:10.1
Unbinding 0000:88:00.1 from ixgbe
Binding 0000:88:00.1 to pciback

[root@meads ~]# /etc/pciback/pciback.sh 0000:98:10.0
Unbinding 0000:98:00.0 from ixgbe
Binding 0000:98:00.0 to pciback

[root@meads ~]# /etc/pciback/pciback.sh 0000:98:10.1
Unbinding 0000:98:00.1 from ixgbe
Binding 0000:98:00.1 to pciback
```

12. View the list of VMs.

```
[root@meads ~]# xm list
Name                                        ID   Mem VCPUs      State
   Time(s)
0004fb0000060000f9b493a2c24f9549             7  8067    16     -b----
   80716.4
Domain-0                                     0  1775    20     r-----
   30379.2
```

13. View the list of assignable devices.

```
[root@meads ~]# xm pci-list-assignable-devices
0000:88:10.0
0000:88:10.1
0000:98:10.0
0000:98:10.1
```

14. Assign these devices to the VM.

Note: In the following example the VM ID is 7.

```
[root@meads ~]# xm pci-attach 7 0000:88:10.0
[root@meads ~]# xm pci-attach 7 0000:88:10.1
[root@meads ~]# xm pci-attach 7 0000:98:10.0
[root@meads ~]# xm pci-attach 7 0000:98:10.1
```

15. View the list of devices for this VM.

```
[root@Meads ~]# xm pci-list 7
Vdev Device
04.0 0000:88:10.0
05.0 0000:88:10.1
06.0 0000:98:10.0
07.0 0000:98:10.1
```

Now VFs on ethernets 9-12 are assigned to VM ID 7, but there are still VFs available to the host. These interfaces appear when you run the **ifconfig** command.

16. Access and log into the OS-E CLI and execute the following.

```
NNOS-E>echo “1” > /sys/bus//pci/rescan
NNOS-E>run ./install_build_mactab.sh
NNOS-E>restart warm
```

After the restart has completed, these interfaces are available on the OS-E.

PCI Passthrough is a specification that allows you to directly connect one VM to one physical device, making the device unavailable to other VMs.

To configure PCI Passthrough:

1. Access and log into the OVM Server’s CLI.

2. Install the necessary packages on the OVM Server.

```
libibumad-1.3.8-2.mlnx1.5.5r2.el5.x86_64.rpm
libibmad-1.3.9-7.mlnx1.5.5r2.el5.x86_64.rpm
opensm-libs-3.3.15-6.mlnx1.5.5r2.el5.x86_64.rpm
kernel-ib-1.5.5.092-2.6.39_300.29.1.el5uek.x86_64.rpm
infiniband-diags-1.5.13.MLNX_20120708-4.mlnx1.5.5r2.el5.x86_64.rpm
ovsvf-config-1.0-6.noarch.rpm
```

3. Create a python script called **vnfs.py** to view and marry PCI addresses to interfaces.

```
#!/usr/bin/python
# Copyright (C) 2012 Steve Jordahl
#
# This program is free software; you can redistribute it and/or modify
```

```
# it under the terms of the GNU Lesser General Public License as
   published
# by the Free Software Foundation; version 2.1 only.
#
# This program is distributed in the hope that it will be useful,
# but WITHOUT ANY WARRANTY; without even the implied warranty of
# MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the
# GNU Lesser General Public License for more details.
#
# vfns: list SR-IOV virtual functions

import os

info = {}

def catFile(filename):
        readfile = open(filename)
        return readfile.read().strip()

for dev in os.listdir('/sys/class/net'):
        if dev.startswith('eth'):
                info[dev] = {}
               info[dev]['address'] = catFile('/sys/class/net/' + dev
   + '/address')

for dev in info.keys():
        devLink = os.readlink('/sys/class/net/' + dev + '/device')
        info[dev]['pci address'] = devLink[-7:]
        os.chdir('/sys/class/net/' + dev)
        for devInfo in os.listdir(devLink):
if devInfo.startswith('virtfn'):
                        info[dev][devInfo] =
   os.readlink(os.path.join(devLink, devInfo))[-7:]

for dev in sorted(info.keys()):
        print dev
        for detail in sorted(info[dev].keys()):
                print "     " + detail + ":  " + info[dev][detail]
```

4. Create **/etc/pciback/pciback.sh**.

```
#!/bin/sh
if [ $# -eq 0 ] ; then
echo "Require a PCI device as parameter"
exit 1
fi
for pcidev in $@ ; do
if [ -h /sys/bus/pci/devices/"$pcidev"/driver ] ; then
echo "Unbinding $pcidev from" $(basename $(readlink /sys/bus/pci/
   devices/"$pcidev"/driver))
echo -n "$pcidev" > /sys/bus/pci/devices/"$pcidev"/driver/unbind
fi
```

```
echo "Binding $pcidev to pciback"
echo -n "$pcidev" > /sys/bus/pci/drivers/pciback/new_slot
echo -n "$pcidev" > /sys/bus/pci/drivers/pciback/bind
done
```

5. Use an IOMMU to allow both the VM and physical devices access to memory. The IOMMU allows the OVM to limit what memory a device is allowed and gives the device the same virtualized memory layout that the guest sees.

```
Edit /boot/grub/grub.conf to enable iommu and comment out the existing
   kernel entry ( see example )

# grub.conf generated by anaconda
#
# Note that you do not have to rerun grub after making changes to this
   file
# NOTICE:  You have a /boot partition.  This means that
#          all kernel and initrd paths are relative to /boot/, eg.
#          root (hd0,0)
#          kernel /vmlinuz-version ro root=/dev/sdb2
#          initrd /initrd-[generic-]version.img
#boot=/dev/sdb
default=0
timeout=5
splashimage=(hd0,0)/grub/splash.xpm.gz
hiddenmenu
title Oracle VM Server-ovs (xen-4.3.0 3.8.13-26.4.2.el6uek.x86_64)
root (hd0,0)
        #kernel /xen.gz console=com1,vga com1=57600,8n1
   dom0_mem=max:1776M allowsuperpage dom0_vcpus_pin dom0_max_vcpus=20
        kernel /xen.gz console=com1,vga com1=57600,8n1
   dom0_mem=max:1776M allowsuperpage
   iommu=passthrough,no-qinval,no-intremap
        module /vmlinuz-3.8.13-26.4.2.el6uek.x86_64 ro
   root=UUID=e2b44279-55a5-48b9-b910-82446b7b8c65 rd_NO_LUKS
   rd_NO_LVM LANG=en_US.UTF-8 rd_NO_MD SYSFONT=latarcyrheb-sun16
   KEYBOARDTYPE=pc KEYTABLE=us rd_NO_DM rhgb quiet
        module /initramfs-3.8.13-26.4.2.el6uek.x86_64.img
```

6. Reboot the OVM Server.

7. Run the vnfs script to view addresses and VFs statistics.

```
[root@meads ~]# ./vnfs.py
eth0
     address:  a0:36:9f:2c:39:74
     pci address:  30:00.0
eth1
     address:  a0:36:9f:2c:39:75
     pci address:  30:00.1
eth10
     address:  00:21:28:a1:e2:41
```

```
     pci address:  88:00.1
eth11
     address:  00:21:28:a1:e2:42
     pci address:  98:00.0
eth12
     address:  00:21:28:a1:e2:43
     pci address:  98:00.1
eth2
     address:  a0:36:9f:2c:39:76
     pci address:  30:00.2
eth3
     address:  a0:36:9f:2c:39:77
     pci address:  30:00.3
eth4
     address:  a0:36:9f:2d:0b:a8
     pci address:  a0:00.0
eth5
     address:  a0:36:9f:2d:0b:a9
     pci address:  a0:00.1
eth6
     address:  a0:36:9f:2d:0b:aa
pci address:  a0:00.2
eth7
     address:  a0:36:9f:2d:0b:ab
     pci address:  a0:00.3
eth8
     address:  00:21:28:a1:e2:46
     pci address:  5f:00.0
eth9
     address:  00:21:28:a1:e2:40
     pci address:  88:00.0
```

8. Run the module.

```
[root@meads ~]# modprobe xen-pciback
```

9. Assign devices to pciback in the format:

```
Domain 0:Bus:#:Device#:Function #).
```

Note: In the following example the 4 interfaces are VFs on ethernets 9-12.

```
[root@meads ~]# /etc/pciback/pciback.sh 0000:88:00.0
Unbinding 0000:88:00.0 from ixgbe
Binding 0000:88:00.0 to pciback

[root@meads ~]# /etc/pciback/pciback.sh 0000:88:00.1
Unbinding 0000:88:00.1 from ixgbe
Binding 0000:88:00.1 to pciback

[root@meads ~]# /etc/pciback/pciback.sh 0000:98:00.0
Unbinding 0000:98:00.0 from ixgbe
Binding 0000:98:00.0 to pciback
```

```
[root@meads ~]# /etc/pciback/pciback.sh 0000:98:00.1
Unbinding 0000:98:00.1 from ixgbe
Binding 0000:98:00.1 to pciback
```

10. View the list of VMs.

```
[root@meads ~]# xm list
Name                                        ID   Mem VCPUs      State
   Time(s)
0004fb0000060000f9b493a2c24f9549             7  8067    16     -b----
   80716.4
Domain-0                                     0  1775    20     r-----
   30379.2
```

11. View the list of assignable devices.

```
[root@meads ~]# xm pci-list-assignable-devices
0000:88:00.0
0000:88:00.1
0000:98:00.0
0000:98:00.1
```

12. Assign these devices to the VM.

Note: In the following example the VM ID is 7.

```
[root@meads ~]# xm pci-attach 7 0000:88:00.0
[root@meads ~]# xm pci-attach 7 0000:88:00.1
[root@meads ~]# xm pci-attach 7 0000:98:00.0
[root@meads ~]# xm pci-attach 7 0000:98:00.1
```

13. View the list of devices for this VM.

```
[root@Meads ~]# xm pci-list 7
Vdev Device
04.0 0000:88:00.0
05.0 0000:88:00.1
06.0 0000:98:00.0
07.0 0000:98:00.1
```

Now ethernets 9-12 are assigned to VM ID 7 and are not available to host any other VMs. They also do not show up when you run the **ifconfig** command.

14. Access and log into the OS-E CLI and execute the following.

```
NNOS-E>echo “1” > /sys/bus//pci/rescan
NNOS-E>run ./install_build_mactab.sh
NNOS-E>restart warm
```

After the restart has completed these interfaces are available on the OS-E.

Installing the OS-E On a VMware ESXi

The OS-E is certified to run on the VMware ESXi 5.5.

Oracle recommends the following configuration.

- vCPUs: 16 (16 sockets, 1 core per socket)
- RAM: 8GB
- Disk: 50G

To install the OS-E on a VMware ESXi:

1. Copy the OS-E's ISO file to the Datastore.
2. Click **Inventory**.
3. Create a new VM by clicking the ESXi server on the left.
4. Select **File > New > Virtual Machine** from the menu.
    - **Configuration**: Select **typical** to accept the default number of CPUs and amount of memory (1 CPU and 1GB). Select **custom** to change the default values. Click **Next**.
    - **Name and Location**: Enter a name for the VM. Click **Next**.
    - **Storage**: Select the Datastore. Click **Next**.
    - **Virtual Machine Version**: *For custom configuration only.* Select **Virtual Machine Version: 8**. Click **Next**.
    - **Guest Operating System**: Select **Linux** for **OS** and **Linux Oracle Linux 4/5/6 (64-bit)** for **Version**. Click **Next**.
    - **CPUs**: *For custom configuration only.* Select the number of sockets and number of cores/sockets. Click **Next**.
    - **Memory**: *For custom configuration only.* Select the memory size. Note the minimum, maximum, and recommended sizes for the guest OS you are using. Click **Next**.
    - **Network**: Select **3. Data Network**. Click **Next**.
    - **SCSI Controller**: *For custom configuration only.* Select **LSI Logic Parallel** (default). Click **Next**.

- **Select a Disk**: *For custom configuration only.* Select **Create a new virtual disk**. Click **Next**.
- **Create a Disk**: Specify the GB for disk capacity and choose **Thick Provision Lazy Zeroed** and **Store with the virtual machine**. Click **Next**.
- **Advanced Options**: *For custom configuration only.* Check the checkbox for **SCSI (0:0)**. Ensure the **Independent** checkbox remains unchecked. Click **Next**.
- **Ready to Complete**: Click **Finish**.

5. Right-click on the VM and select **Edit Settings...**.
   - Select **CD/DVD Drive 1**.
   - **Device Status**: Select **Connect at power on**.
   - **Device Type**: Select **Datastore ISO File** and choose **install_<*release_version*>_<*build_number*>.iso on Datastore1**.
   - Click **OK**.
6. Power on the VM by clicking the green play button.
7. Right-click the VM and select **Open Console**.

The OS-E is now ready to be set up and configured.

Configuring ESXi Passthrough

On the ESXi, you can directly connect a VM to a physical port via the SR-IOV specification. SR-IOV treats a single physical device as multiple separate Virtual Functions (VF)s. To deploy SR-IOV, you must enable VFs at the host level.

To configure SR-IOV on the ESXi you must have a NIC with an intel 82599 chipset or newer and a BIOS, both supporting SR-IOV.

The configuration for SR-IOV on the ESXi consists of two parts: first you must configure the OS-E's VM server, then you must assign individual VFs to specific VMs.

To configure the OS-E's VM server for SR-IOV:

1. Enable SR-IOV in the BIOS.

2. Ensure you have the latest drivers for your intel NIC (ixgbe) and ESXi version. See https://my.vmware.com/web/vmware/info/slug/datacenter_cloud_infrastructure/vmware_vsphere_with_operations_management/5_5#drivers_tools for more information on ESXi 5.5 drivers.
3. Install the appropriate drivers and reboot the host.
4. Log into the ESXi CLI shell and enter the following command to view a list of all NICs on the server and identify which NICs to use for SR-IOV.

```
# lspci | grep -i 'ethernet\|network'
```

5. Specify the number of VFs you are assigning to each port by executing the following command:

```
# esxcfg-module ixgbe -s max_vfs=<P1=n><P2=n><P3=n><P4=n>
```

where <*Px*=*n*> stands for the configured ports and their assigned VF values, less than or equal to 63. Assigning a value of 0 makes that port unavailable for SR-IOV.

Note: The SR-IOV specification allows for you to partition the Physical Function (PF) into a particular number of VFs you can then attach to VMs. The maximum number of VFs you can create on a PF depends on the hardware you are using. Typically, for 10GbE chipsets equal to or newer than 82599, that number is 63.

6. Verify that you entered the correct values by entering the following command:

```
# esxcfg-module ixgbe -g ixgbe
```

7. Reboot the server.
8. View the list of configured VFs by either reentering the following command:

```
#lspci | grep -i 'ethernet\|network'
```

or accessing, via the vSphere GUI, **Host** > **Configuration** > **Advanced Settings**.

To configure a specific OS-E VM for SR-IOV:

Note: To attach a VF to a VM, the VM version must be greater than or equal to 10.

1. Power off the VM.
2. Select **Settings** > **Hardware** > **Add**.
3. Select **PCI device** and select the VF you are adding to the VM.
4. Repeat this procedure for each VF you are adding to the VM.

Note: If you are prompted to “reserve” resources, you may have to click that button for the VM to power on.

Once a VF is attached to a particular VM, you cannot attach it to any other VM.

Installing the OS-E As a XEN Virtual Machine

The OS-E is certified to run on XEN 3.4.3.

Oracle recommends the following configuration.

- vCPUs: 16 (16 sockets, 1 core per socket)
- RAM: 8GB
- Disk: 50G

  Note: Oracle recommends using LVM partitions as disks.

1. Create a partition; the following example creates a 50G partition.

   Note: The following example assumes the volume group is named “ol”.

```
# lvcreate --size=50G --name=asc ol
```

   Logical volume “asc” is created.

2. Download OL7.2 ISO image from the Oracle Software Delivery Cloud and copy the file to the /tmp directory on the server.
3. Create a config file for the VM at /etc/xen/asc.cfg. The following is an example config file.

   Note: The following is an example. Ensure you customize your config file, including changing the MAC addresses, to fit your environment.

```
#  -*- mode: python; -*-
#============================================================================
# Python configuration setup for 'xm create'.
# This script sets the parameters used when a domain is created using
   'xm create'.
# You use a separate script for each domain you want to create, or
# you can set the parameters for the domain on the xm command line.
#============================================================================

#----------------------------------------------------------------------------
# PV GRUB image file.
kernel = "/usr/lib/xen/boot/hvmloader"
builder = 'hvm'
```

```
device_model = '/usr/lib64/xen/bin/qemu-dm'

# Sets path to menu.lst
extra = "(hd0,1)/grub/menu.lst"
# can be a TFTP-served path (DHCP will automatically be run)
# extra = "(nd)/netboot/menu.lst"
# can be configured automatically by GRUB's DHCP option 150 (see grub
   manual)
# extra = ""

# Initial memory allocation (in megabytes) for the new domain.
#
# WARNING: Creating a domain with insufficient memory may cause out of
#          memory errors. The domain needs enough memory to boot kernel
#          and modules. Allocating less than 32MBs is not recommended.
memory = 8192

# A name for your domain. All domains must have different names.
name = "asc"

# 128-bit UUID for the domain.  The default behavior is to generate a
   new UUID
# on each call to 'xm create'.
#uuid = "06ed00fe-1162-4fc4-b5d8-11993ee4a8b9"

# List of which CPUS this domain is allowed to use, default Xen picks
#cpus = ""         # leave to Xen to pick
#cpus = "0"        # all vcpus run on CPU0
#cpus = "0-3,5,^1" # all vcpus run on cpus 0,2,3,5
#cpus = ["2", "3"] # VCPU0 runs on CPU2, VCPU1 runs on CPU3

# Number of Virtual CPUS to use, default is 1
vcpus = 4
cpus = "4-31" # all vcpus run on cpus >3

#---------------------------------------------------------------------
# Define network interfaces.

# By default, no network interfaces are configured.  You may have one
   created
# with sensible defaults using an empty vif clause:
#
# vif = [ '' ]
#
# or optionally override backend, bridge, ip, mac, script, type, or
   vifname:
#
# vif = [ 'mac=00:16:3e:00:00:11, bridge=xenbr0' ]
#
# or more than one interface may be configured:
#
```

```
# vif = [ '', 'bridge=xenbr1' ]

vif = [ 'mac=00:16:3E:80:00:18, bridge=management',
    'mac=00:16:3E:80:00:19, bridge=data', 'mac=00:16:3E:80:00:1A,
    bridge=messaging' ]

#----------------------------------------------------------------------
# Define the disk devices you want the domain to have access to, and
# what you want them accessible as.
# Each disk entry is of the form phy:UNAME,DEV,MODE
# where UNAME is the device, DEV is the device name the domain will
    see,
# and MODE is r for read-only, w for read-write.

disk = [ 'phy:/dev/mapper/ol-asc,hda,w', 'file:/tmp/
    V100082-01-OL7U2-x86_64.iso,hdc:cdrom,r']
#----------------------------------------------------------------------
# Define frame buffer device.
#
# By default, no frame buffer device is configured.
#
# To create one using the SDL backend and sensible defaults:
#
# vfb = [ 'sdl=1' ]
#
# This uses environment variables XAUTHORITY and DISPLAY.  You
# can override that:
#
# vfb = [ 'sdl=1,xauthority=/home/bozo/.Xauthority,display=:1' ]
#
# To create one using the VNC backend and sensible defaults:
#
# vfb = [ 'vnc=1' ]
#
# The backend listens on 127.0.0.1 port 5900+N by default, where N is
# the domain ID.  You can override both address and N:
#
# vfb = [ 'vnc=1,vnclisten=127.0.0.1,vncdisplay=1' ]
#
# Or you can bind the first unused port above 5900:
#
# vfb = [ 'vnc=1,vnclisten=0.0.0.0,vncunused=1' ]
#
# You can override the password:
#
# vfb = [ 'vnc=1,vncpasswd=MYPASSWD' ]
#
# Empty password disables authentication.  Defaults to the vncpasswd
# configured in xend-config.sxp.

#----------------------------------------------------------------------
```

```
# Define to which TPM instance the user domain should communicate.
# The vtpm entry is of the form 'instance=INSTANCE,backend=DOM'
# where INSTANCE indicates the instance number of the TPM the VM
# should be talking to and DOM provides the domain where the backend
# is located.
# Note that no two virtual machines should try to connect to the same
# TPM instance. The handling of all TPM instances does require
# some management effort in so far that VM configration files (and thus
# a VM) should be associated with a TPM instance throughout the
   lifetime
# of the VM / VM configuration file. The instance number must be
# greater or equal to 1.
#vtpm = [ 'instance=1,backend=0' ]

#----------------------------------------------------------------------
# Configure the behaviour when a domain exits.  There are three
   'reasons'
# for a domain to stop: poweroff, reboot, and crash.  For each of these
   you
# may specify:
#
#   "destroy",        meaning that the domain is cleaned up as normal;
#   "restart",        meaning that a new domain is started in place of
   the old
#                     one;
#   "preserve",       meaning that no clean-up is done until the domain
   is
#                     manually destroyed (using xm destroy, for
   example); or
#   "rename-restart", meaning that the old domain is not cleaned up,
   but is
#                     renamed and a new domain started in its place.
#
# In the event a domain stops due to a crash, you have the additional
   options:
#
#   "coredump-destroy", meaning dump the crashed domain's core and then
   destroy;
#   "coredump-restart', meaning dump the crashed domain's core and the
   restart.
#
# The default is
#
#   on_poweroff = 'destroy'
#   on_reboot   = 'restart'
#   on_crash    = 'restart'
#
# For backwards compatibility we also support the deprecated option
   restart
#
# restart = 'onreboot' means on_poweroff = 'destroy'
```

```
#                             on_reboot   = 'restart'
#                             on_crash    = 'destroy'
#
# restart = 'always'   means on_poweroff = 'restart'
#                             on_reboot   = 'restart'
#                             on_crash    = 'restart'
#
# restart = 'never'    means on_poweroff = 'destroy'
#                             on_reboot   = 'destroy'
#                             on_crash    = 'destroy'

#on_poweroff = 'destroy'
#on_reboot   = 'restart'
#on_crash    = 'restart'
#---------------------------------------------------------------------
#   Configure PVSCSI devices:
#
#vscsi=[ 'PDEV, VDEV' ]
#
#   PDEV   gives physical SCSI device to be attached to specified guest
#          domain by one of the following identifier format.
#          - XX:XX:XX:XX (4-tuples with decimal notation which shows
#                         "host:channel:target:lun")
#          - /dev/sdxx or sdx
#          - /dev/stxx or stx
#          - /dev/sgxx or sgx
#          - result of 'scsi_id -gu -s'.
#            ex. # scsi_id -gu -s /block/sdb
#                  36000b5d0006a0000006a0257004c0000
#
#   VDEV   gives virtual SCSI device by 4-tuples (XX:XX:XX:XX) as
#          which the specified guest domain recognize.
#

#vscsi = [ '/dev/sdx, 0:0:0:0' ]

#=====================================================================

# Guest VGA console configuration, either SDL or VNC
#sdl = 1
vnc = 1
vncpasswd=""
vncdisplay=10
vnclisten="0.0.0.0"
```

4. Start the VM.

```
# xl start /etc/xen/asc.cfg
```

5. Vnc to host:10 to start the OL7 installation process.

Once OL7 is installed, you can begin the OS-E installation. For more information on installing OS-E software, see Chapter 2, Installing the OS-E System.

Installing the OS-E On KVM

The OS-E is certified to run on KVM on OL7.

Oracle recommends the following configuration.

- vCPUs: 8
- RAM: 8GB
- Disk: 50G

  Note: Oracle recommends using LVM partitions as disks.

1. Install the KVM packages.

```
# yum install kvm libyirt
# yum install python-virtinst virt-top virt-manager virt-v2v
   virt-viewer
```

2. Use the virt-manager command to create your networks.
3. Install the OS-E guest by either using the following command in the CLI or via the virt-manager (right-click localhost (QEMU) and click New).

```
virt-install -n asc -r 8192 --os-type=linux --disk /dev/mapper/
   ol-asc,device=disk,bus=virtio,size=50,sparse=false,format=raw -w
   network=management,model=virtio -w network=messaging,model=virtio
   -w network=data,model=virtio -c /mnt/install/<build_version>.iso
   --vcpus=8
```

Configuring the VM

Once the VM is installed and running, you now must configure it to match the SIP application you are supporting. Since the VM does not have a pre-installed base configuration, Oracle provides the **config setup** configuration setup script that you can use to create a base configuration.

Using Config Setup

For Oracle users who are familiar with OS-E, the *config setup* script enables the configuration on the VM to make it reachable via ICMP (ping), SSH, and HTTPS for further configuration. The script presents a set of questions to help you with the initial system configuration. The information in the script includes the following:

- Local hostname
- IP interface names and addresses
- SSH and Web access
- Default route and any additional static routes per interface for remote management
- User-defined OS-E

Every Oracle OS-E system has a minimum of two Ethernet interfaces. Any Ethernet interface on the system can be used for management traffic, however, Oracle recommends the use of eth1, as eth0 is reserved for fault-tolerant clustering with other OS-E systems. Management traffic is also supported on any interface that is carrying private or public network traffic. This means that it would be possible to use eth1 to carry SIP traffic and management traffic.

CLI Session

```
NNOS-E-VM> config setup
set box\hostname: <name>
config box\interface: eth1
set box\interface eth1\ip a\ip-address: <ipAddress/mask>
config box\interface eth1\ip a\ssh (y or n)? n
config box\interface eth1\ip a\web (y or n)? y
config box\interface eth1\ip a\routing\route: <routeName>
set box\interface eth1\ip a\routing\route localGateway\gateway:
<ipAddress>
set box\cli\prompt: <newPrompt>
Do you want to commit this setup script (y or n) y
Do you want to update the startup configuration (y or n)? y
```

Sample VM Configuration

This section describes a base configuration designed to support a standard SBC application where the VM functions with SIP endpoints and a PBX or feature server. The high-level details of this configuration are provided below and additional details are embedded in the configuration file itself at the end of this section.

- Two interfaces: one "outside" and one "inside."
- Management ports for ICMP, SSH, and HTTPS open on both interfaces.
- The IP address associated with a DNS resolver.
- SIP UDP, TCP, and TLS ports open on both interfaces.
- NAT traversal & media anchoring enabled.
- A sample gateway configuration for an attached PBX or feature server.
- A sample registration- and dial-plan for delegation of SIP traffic to the attached PBX or feature server.
- A local registration plan to support registrations and calls locally through the VM (for cases where there is no attached PBX or feature server).

Note: Oracle recognizes that the items in the base configuration will not be 100% applicable to all OS-E-VM deployments. However, by including these items in this sample configuration, new VM users can observe the configuration structure and hierarchy. Any necessary changes to this base configuration can be made using the procedures described in the Oracle manual set. See, “Using Oracle Documentation,” for more information.

Below is a copy of the base configuration. Note that any changes to the configuration should be made using the OS-E Management System (see, “Enabling the OS-E Management System”).

Note: Oracle does not recommend editing the configuration file below directly, and then importing it into the VM. While the VM does support this function, it is possible to introduce syntax errors into the configuration file using this method. Modifying the configuration with the CLI or Management System prevents this possibility.

This section is unique to every VM; you do not need to edit this.

```
config cluster
 set name acmepacket-nnos-e-vm-demo
 config box 1
  set hostname acmepacket-nnos-e-vm-demo
  set name acmepacket-nnos-e-vm-demo
  set identifier 00:0c:29:c9:7a:e2
```

The IP address is configured as part of the configuration script execution.

```
config interface eth0
   config ip outside
    set ip-address static 172.30.3.128/22
    config ssh
    return
    config web
    return
    config sip
     set nat-translation enabled
     set udp-port 5060 "" "" any 0
     set tcp-port 5060 "" "" any 0
     set tls-port 5061 "" "" any 0
     set certificate vsp\tls\certificate sample
    return
    config icmp
    return
    config media-ports
    return
    config routing
     config route default
      set gateway 172.30.0.1
    return
    return
   return
  return
```

The following section of the configuration provides a DNS resolver entry and is configured as part of the configuration script execution. This is not required for operation but can be helpful if you want to use FQDNs in the config instead of IPs)

```
config dns
  config resolver
   set server 192.168.1.3 UDP 53 100 ALL
  return
 return
return
```

The following IP is disabled; you can enable it once you change the IP to match your local network conditions.

```
config interface eth1
   config ip inside
```

```
    set admin disabled
    set ip-address static 192.168.1.2/24
    config ssh
    return
    config web
    return
    config sip
     set udp-port 5060 "" "" any 0
     set tcp-port 5060 "" "" any 0
     set tls-port 5061 "" "" any 0
     set certificate vsp\tls\certificate sample
    return
    config icmp
    return
    config media-ports
    return
```

This routing config is provided as an example; edit it as needed. Change to match your preferred NTP server.

```
config routing
     config route inside-ntwk
      set destination network 192.168.0.0/16
      set gateway 192.168.1.1
   return
    return
   return
  return
  config ntp-client
   set server pool.ntp.org
return
  config cli
   set prompt nnos-e-vm>
  return
 return
return
```

The following section of the configuration contains all of the event log filters and targets.

```
config services
 config event-log
  config file eventlog
   set filter all error
  return
  config file access-log
   set filter access info
  return
  config file kernelsys
   set filter krnlsys debug
  return
```

```
  config file db
   set filter db debug
  return
  config file system
   set filter general info
   set filter system info
  return
  config file access
   set filter access info
  return
  config file dos
   set filter dosSip alert
  return
  config local-database
   set filter all error
  return
 return
return
```

The following section of the config provides some commonly used default system parameters; further information on these parameters is provided in the tech manuals.

```
config master-services
 config database
  set media enabled
 return
return

config vsp
 set admin enabled
 config default-session-config
  config media
   set anchor enabled
   config nat-traversal
    set symmetricRTP true
   return
   set rtp-stats enabled
  return
  config sip-directive
   set directive allow
  return
  config log-alert
  return
 return
 config tls
  config certificate sample
  return
 return
```

The following section of the configuration provides a sample policy rule to reject calls from a user with a URI that starts with 1000. This sample is provided as a means of introducing a new user to the concept of policy rules)

```
config policies
  config session-policies
   set default-policy vsp\policies\session-policies\policy default
   config policy default
    config rule sample-rule
     set description "sample rule to reject calls"
     config condition-list
      set from-uri-condition user match 1000
     return
     config session-config
      config sip-directive
       set directive refuse 400 "Please Pay Your Bill"
      return
     return
    return
   return
  return
 return
```

The folllowing configuration provides a sample dial-plan that takes a call with a Req URI domain of delegate.com and forwards it to the sample SIP gateway.

```
 config dial-plan
  config route sample-delegate
   set description "delegate to defined server"
   set peer server "vsp\enterprise\servers\sip-gateway sample-gateway"
   set request-uri-match domain-exact delegate.com
  return
 return
```

The following configuration provides a sample registration plan that takes a registration attempt with a domain of *xyz.com* and registers the endpoint locally. This is useful for cases where you want to register an endpoint locally for call testing purposes.

```
config registration-plan
  config route sample-accept-local
   set description "accept registers locally for this domain"
   set action accept
   set peer server "vsp\enterprise\servers\sip-gateway sample-gateway"
   set to-uri-match domain-exact xyz.com
  return
```

The following configuration provides a sample registration plan that takes a registration attempt with a domain of delegate.com and proxies the registration to the attached PBX or feature server.

```
config route sample-delegate
   set description "delegate to the defined server"
   set peer server "vsp\enterprise\servers\sip-gateway sample-gateway"
   set to-uri-match domain-exact delegate.com
  return
 return
```

The following configuration provides a sample SIP gateway that could be used for an attached PBX or feature server. You will need to edit the IP address to reflect the actual server IP or FQDN.

```
 config enterprise
  config servers
   config sip-gateway sample-gateway
    config server-pool
     config server sample-server
      set host 192.168.1.4
     return
    return
   return
  return
 return

config external-services
return
config preferences
 config cms-preferences
 return
return
```

The following configuration provides two different sample permission sets. These permission sets modified and/or can be used with user accounts that you create.

```
config access
 config permissions super-user
  set cli advanced
 return
 config permissions view-only
  set cli disabled
  set ftp disabled
  set config view
  set actions disabled
  set templates disabled
  set web-services disabled
  set debug disabled
 return
```

```
return

config features
return
```

Acme Packet recommends that the storage-device fail-threshold be set to 200 MB.

```
services
storage-device
  fail-threshold 200 MB
```

Enabling the OS-E Management System

Once you have configured an Ethernet interface, such as eth1, you can use your Internet Explorer Web browser to point to the configured IP address of this interface to launch the OS-E Management System. The OS-E Management System provides a windows and menu user interface to configuring the OS-E. See the *Net-Net OS-E – Using the NNOS-E Management Tools* for information on using the CMS.

Bridging to Additional Ethernet Ports

Follow the steps in this section if you need to configure VMware Player on a Window platform to use two bridged networks. By default, VMWare Player allows the following functionality:

- One bridged interface (to the first host network interface)
- One NAT interface
- One host-only interface

To create two bridged interfaces, you will need to

1. add an additional VMnet associated with a second interface, and
2. edit the VM configuration file to use the new VMnet.

Adding an Additional VMnet

To add an additional VMnet, perform the following steps:

1. Halt all VMs currently running on this x86-based PC or server.

2. Launch the **vmnetcfg.exe** application from the VMware Player installation directory (c:\Program Files\VMware\VMware Player\vmnetcfg.exe).
3. Select the **Host Virtual Network Mapping** tab.
4. Select a VMnet to use for the second network interface card (NIC), such as VMnet3.
5. From the drop-down men, select the NIC you wish to connect to this VMnet.

If you want to have more control over which VMnet0 which connects to the first NIC perform the following steps:

1. Select the **Automatic Bridging** tab.
2. In the **Automatic Bridging** box, de-select the **Automatically choose and available physical network adapter to bridge to VMnet0**.
3. Select the **Host Virtual Network Mapping** tab.
4. Select a VMnet to use for the first NIC, such as VMnet2.
5. From the drop-down menu, select the NIC you wish to connect to this VMnet.

Note: You can use VMnet0 to assign to a specific NIC. However, avoiding VMnet0 will indicate to a later user of the VMs configuration file that specific NICs were assigned to the VMs virtual interfaces, thus removing any questions about the automatic nature implied with VMnet0 on any particular system.

Editing the VM Configuration File

You will need to edit the VMware configuration file to include the second NIC with the VMware Player. Perform the following steps:

1. Halt all VMs currently running on this x86-based PC or server.
2. Using Windows Explorer, open the Oracle OS-E folder.

- Using a text editor such as Notepad, open the file **nnos-e-vm.vmx**.
- At the bottom of the file add the following lines, substituting the desired VMnets for the Ethernet interfaces:
  - **ethernet0.connectionType = "custom"**
  - **ethernet0.vnet = "vmnet0"**

  - **ethernet1.connectionType = "custom"**
  - **ethernet1.vnet = "vmnet3"**
- Ensure that there are no other lines in the file specifying **ethernet<u>X</u>.connectionType = "<u>XXXXX</u>"**.

OS-E-VM Troubleshooting

Oracle makes every effort to test the VM in a variety of customer environments. This section covers recently reported issues directly from OS-E-VM customers. If you discover an issue with the VM that we need to know about, contact Oracle Customer Support for assistance.

Installing the VM on Slow Systems

There may be cases where you are allowed to log in to the VM before the Oracle OS-E application starts up. This is usually caused by a slow system where you installed the VM, and all necessary software processes are requiring more time to complete their startup routines.

If you shut down the VM, and then start it again later, the VM will return you to the same screen and prompt that was displayed at the time of the shutdown. The VM startup would appear as shown in the following image.

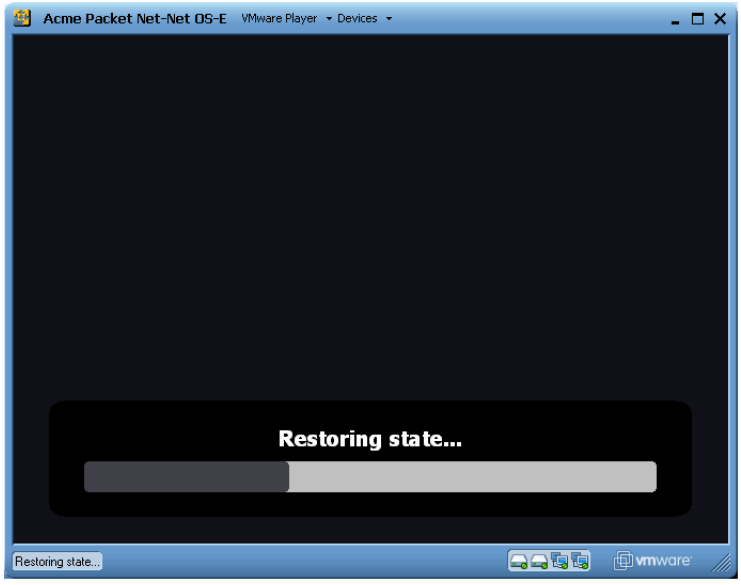

When the VMWare system returns to the state where you previously logged out, and if you restart VMWare, you will see the login: prompt, such as the hostname *mikeo-cva* in the following image. If you did not previously configure a unique *username* and *password*, just press **ENTER** at the username and password prompts.

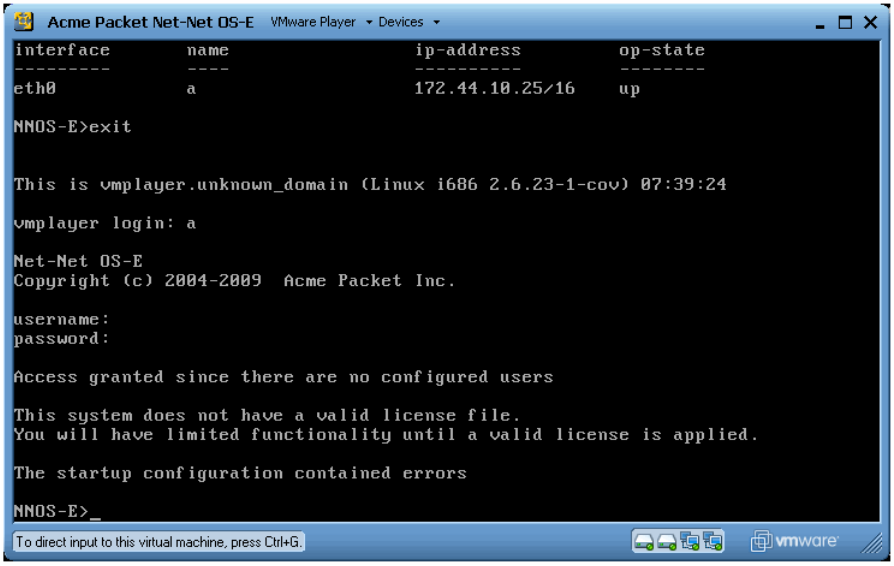

If the Oracle OS-E-VM application has not yet started, your screen appears as shown in the following image. Exit and wait a few minutes for the application to complete all the essential startup processes. You can then log in by pressing ENTER/ ENTER at the username and password prompts, or you can enter the previously-configured username and password.

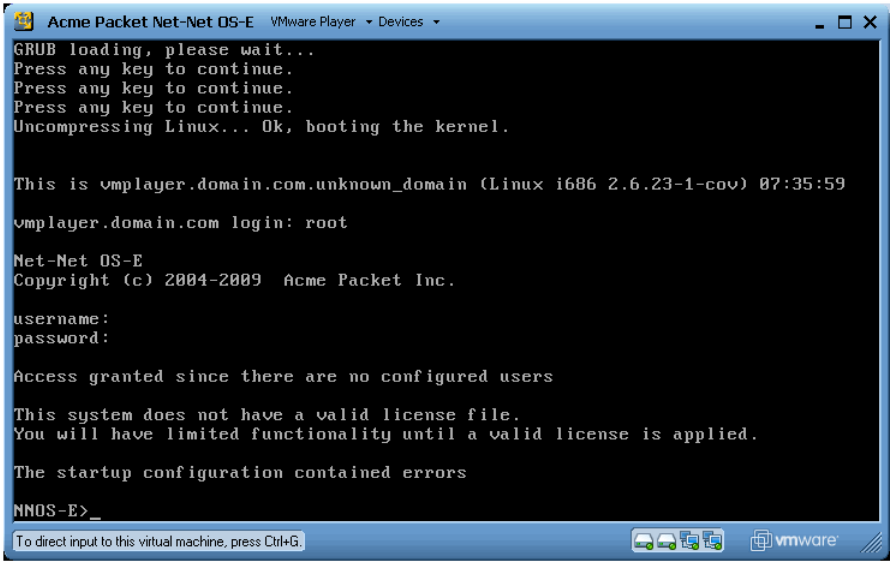

Other VM Limitations and Considerations

The following limitations and considerations apply:

- Configuring feature options that rely on critical timing are more problematic to VMs. This includes music-on-hold (MoH), announcements, periodic announcements, and transcoding.
- VMs running on AMD systems exhibit more timing issues than on Intel-based systems.
- The type of hardware over which you are running the VM can make a significant difference in VM performance. Improved performance is normal when running the VM over larger and faster running platforms.