Explorer Controller Suite 3.0

Installation and Upgrade Guide

Please Read

Important

Read this entire guide. If this guide provides installation or operation instructions, give particular attention to all safety statements included in this guide.

Notices

Trademark Acknowledgments

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: **www.cisco.com/go/trademarks**.

Third party trademarks mentioned are the property of their respective owners.

The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

Publication Disclaimer

Cisco Systems, Inc. assumes no responsibility for errors or omissions that may appear in this publication. We reserve the right to change this publication at any time without notice. This document is not to be construed as conferring by implication, estoppel, or otherwise any license or right under any copyright or patent, whether or not the use of any information in this document employs an invention claimed in any existing or later issued patent.

Copyright

Contents

About This Guide

Introduction

This document contains information regarding the installation and upgrade of the Explorer Control Suite (ECS) to release 3.0.

The ECS provides system operators with a regional command and control interface that manages collocated or geographically dispersed system controllers, such as Cisco's Explorer Controller (EC) and Digital Transport Adapter Control System (DTACS).

The main features of ECS 3.0 are to provide the following functionality as independently deployable services:

- Updated Linux OS platform from RHEL 5.x to CentOS 6.x.
- Improved hardware abstraction & deployment automation; VMWare support.
- Improved failover and Mean Time To Repair (MTTR) due to simplified procedures; elimination of manual procedures associated with RMAN Backup and Restore.
- CPEMS enhancement now permanently deletes DHCT records from the CPEMS database.

Audience

This guide is written for field service engineers and system operators who are responsible for installing, migrating or upgrading to ECS 3.0.

Required Skills and Expertise

System operators or engineers who install or upgrade the ECS software need the following skills:

- Advanced knowledge of UNIX
  - Experience with the UNIX vi editor. Several times throughout the system upgrade process, system files are edited using the UNIX vi editor. The UNIX vi editor is not intuitive. The instructions provided in this guide are no substitute for an advanced working knowledge of vi.
- Knowledge of VMware
- Knowledge of Linux
- Knowledge or PowerShell

Document Version

This is the first formal release of this document.

Revision History

Date	Revision	Section
20170907	■ Updated command in Step 6 to preserve attributes for copied file. ■ Added Step 10 to define access permissions for config.json file.	***Configuring the Consul Configuration File*** (on page 168)
20170919	Deleted "Exclude VCS Console Service and Alarm Manager Services From YUM Updates" section.	***Install the VCS Console*** (on page 54)
20170919	Corrected command syntax and example in Step 7.	***Upgrading the VCS Console Servers*** (on page 196)
20170925	Updated section including "Important" content and procedure.	***Upgrading the BOA Servers*** (on page 200)
	Deleted "Verifying the Enterprise Manager" section in Chapter 2.	N/A
	■ Updated Step 1 (missing s): yum install ServicesManager ■ Corrected term in brackets to: VCS_Console_HA_IP	***Configuring snmpd on Consul Node*** (on page 51) ***Configuring snmpd on the VCS Console Node*** (see on page 65) ***Configuring snmpd on the ECS Node*** (on page 89)
	Added "Configuring snmpd on the BOA Node"	***Configuring snmpd on the BOA Node*** (on page 102)
	Modified procedure in "Import CPE Data From the EC to the ECS" section	***Import CPE Data From the EC to the ECS*** (on page 181)

1

Planning the Install or Migration

Introduction

This chapter contains information that helps you and Cisco engineers plan the installation or migration to minimize system downtime.

In This Chapter

Hardware Requirements

The following hardware prerequisites are required to deploy virtual machines (VMs) in an ECS 3.0 environment.

- The Oracle Real Application Clusters (RAC) and UCS blades are racked, cabled, and configured
- Requires Cisco UCS hardware (C240 M3 or C240 M4) with the latest ESXi software installed
- Requires Admin Node Virtual Machine.
- Cisco UCS hardware should have adequate CPU, Memory, a local disk datastore and a sufficient network for the following virtual machines (VMs).

VM	CPUs	Memory (GB)	Hard Disk (GB)	Network Interfaces	# Nodes for HA
Oracle RAC	8	64	1 x 80 1 x 10 1 x 512 1 x 1536	1 x Public 1 x Private	2
Consul	1	2	1 x 32	1 x Public	3
VCS Console	8	16	1 x 160	1 x Public 1 x Private	2
Alert Manager	2	2	1 x 32	1 x Public	2
ECS	4	4	1 x 32	1 x Public 1 x Private	2
BOA	4	4	1 x 32	1 x Public	2

Software Requirements

The following software prerequisites are required for an ECS 3.0 environment.

- vCenter 5.5 or later, PowerShell 3.0, and PowerCLI 4.10.793510 or later are installed on a PC.
- Requires a vCenter Web UI login or a vSphere client to connect and perform management tasks; vCenter login must have admin privileges to deploy VMs .
- Requires Admin Node access:
  - The Admin Node should already be installed and configured.
  - The Admin Node was built using the ***Admin Node Installation Guide*** (part number TP-00145).
- At least one external Network Time Protocol (NTP) server, version 4.x or later, configured and accessible on the network.
- A vCenter user name and datacenter created for you that include privileges to create hosts and add VMs.
- A naming schema that is unique for your datacenter.
- An NFS storage area for ECS and Oracle-related files.
- ECS 3.0-specific software.

  Important: All software required for the ECS installation should exist on the Admin Node deployed and installed in your NextX system. If it is not yet downloaded to your Admin Node, refer to the ***Admin Node Installation Guide*** for details.

X.509 CA Certificate and Associated Private Key Requirements

Each node in your ECS 3.0 system requires a NextX X.509 certificate along with an associated private key. The X509 certificates must be signed by a Certification Authority (CA). The CA can be either an external entity or an internal CA can be created on the Admin Node.

Important: During the installation and configuration of the Admin Node, you should have created a root CA, as well as all of the certificate/key pairs for each node in your ECS 3.0 system. If you have not created these certificates, refer to Chapters 5 through 6 in the ***Admin Node Installation Guide*** to create them now.

The NextX X.509 certificates were created when you deployed and configured the Admin Node. In this guide, you will distribute the certificates from the Admin Node to their respective node on the ECS system.

2

RAC Installation

Introduction

This chapter describes the installation procedures for installing Oracle Database 11g release 2 (11.2) with Oracle Real Application Clusters (Oracle RAC).

Important: The procedures in this chapter are only for initial installation of ECS 3.0. If you are performing an ECS migration, your RAC system is already built. Go to ***Install and Configure the ECS System*** (on page 43).

In This Chapter

Blade Configuration

Important:

- If all prerequisite tasks have been completed, you should have a vCenter, with an ECS DataCenter, containing at least four hosts (blades).
- If this is a default installation, you will have two chassis, with two blades in each. Your vCenter should be configured with an ECS Datacenter containing all four hosts (blades).
- If this is not a default installation and the Datacenter and hosts are already present, the Oracle RAC installation allows you to select the host where the RAC VMs will be installed.

Detailed instructions for installing the Oracle RAC are in the following procedures.

Example:

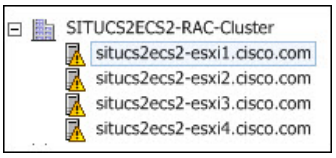

Adding ESXi Hosts to the Datacenter

Important: Complete this procedure only if all four hosts have not been added to the ECS Datacenter.

If the UCS blades have not been added to the Datacenter, complete the following procedure to add them.

1 Right-click the Datacenter and select **Add Host**. The Add Host Wizard window appears.

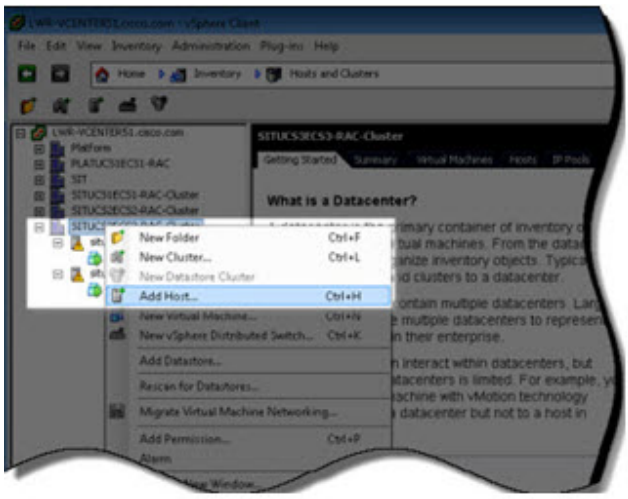

2 In the **Connection** area, enter the following information and then click **Next**. A Security Alert window appears.

- **Host** – The DNS name configured for the ESXi Blade
- **Username** – The ESXi root login (such as, root)
- **Password** – The password of the ESXi root user

Examples:

- **Host**: situcs2ecs2-esxi1
- **Username**: root
- **Password**: The root password

3 Click **Yes** to verity the authenticity of the host. The Host Summary area appears.

4 Review the content of this page, and if the information is correct, click **Next**. The Assign License area appears.

Note: If after clicking **Next** you realize the information was not correct, click the **Back** button to return to the Add a Host window. Correct the entries and click **Next**.

5 Select a License Key and click **Next**. The Lockdown Mode area appears.

Note: You can select any License Key that is listed.

6 Retain the default setting and click **Next**. The Virtual Machine Location area appears.

7 Select the Datacenter and click **Next**. The Ready to Complete area appears.

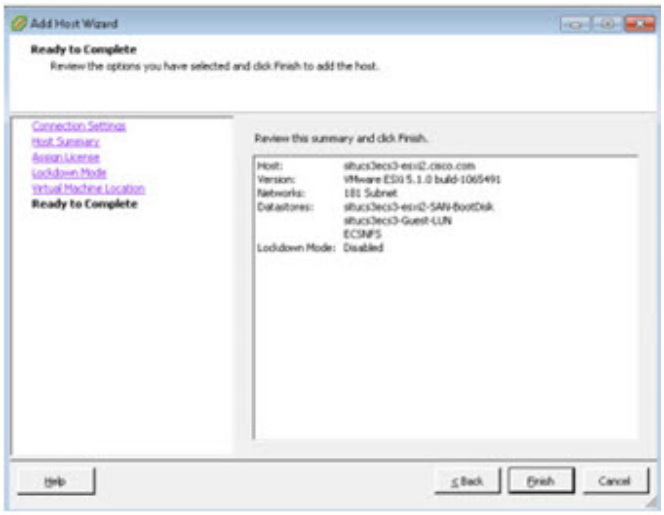

8 Review the content of this page. If the information is correct, click **Finish**.

Note: If the information is not correct, click the **Back** button to go back to the appropriate pages.

9 Verify that ESXi host1 was added to the Datacenter.

10 Repeat this procedure to add ESXi Hosts 2, 3, and 4.

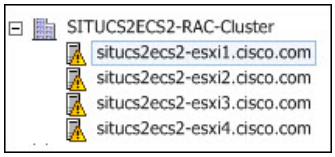

Adding Storage Devices to ESXi Hosts

Complete the following steps to add the required storage devices to each ESXi host. The first time you access Hardware Storage, you will see only one Datastore. This Datastore will be named datastore1. It is actually the SAN-BootDisk.

In this procedure you will rename datastore1 and rescan the device for the other required devices. After completing this procedure, the following storage devices will be present:

- Oradata-LUN1
- Guest-LUN1
- SAN-BootDisk
- Local-Disk1 (Optional)

Important: ESXi hosts 2 and 4 will not have the Oradat-LUN1 datastore.

Note: The Local-Disk1 is not necessary or required.

1. In vCenter, click **ESXi host 1** and then click the **Configuration** tab.
2. In the Hardware box, click **Storage**.
3. In the Datastore window, verify whether or not the SAN-BootDisk and Guest-LUN1 Datastores are present.
   - If they are present, repeat Steps 1 through 3 on the next ESXi host.
   - If they are not present, go to Step 4.
4. From the Datastores window, select **datastore1** (size 15GB).

5 From the Datastore Details section, click **Properties**. The Volume Properties window appears.

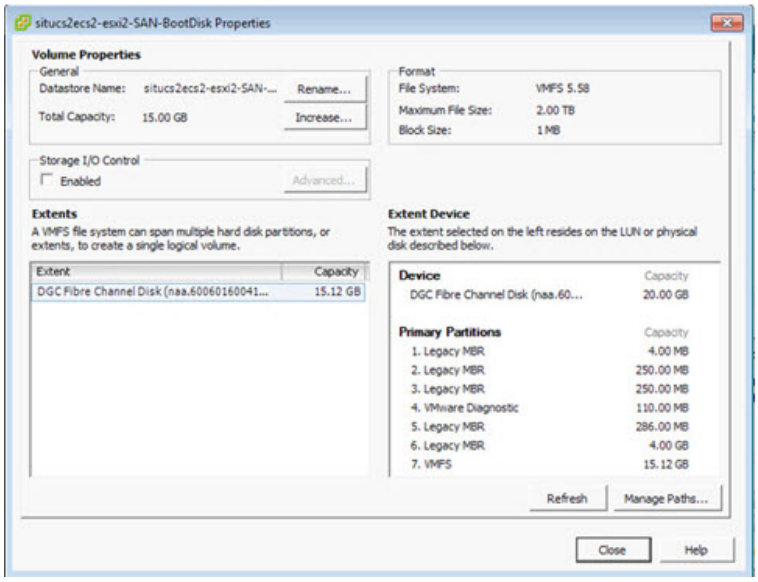

6 In the General area, click **Rename**.

7 Type the name of the Datastore and click **OK**.

Example: SAN-BootDisk

8 Click **Close**.

9 From the Datastores section, click **Rescan All** to rescan the available Datastores. The Rescan window appears.

10 Retain all of the default settings; then, click **OK**.

Important: The system should rescan, find, and then add the Guest-LUN1 Datastore.

Note: You should now have the following Datastores:

- Oradat-LUN1
- Guest-LUN1
- SAN-BootDisk
- Local-Disk1 (Optional)

11 Repeat this procedure to add the storage devices to the remaining ESXi host.

Datastore Example

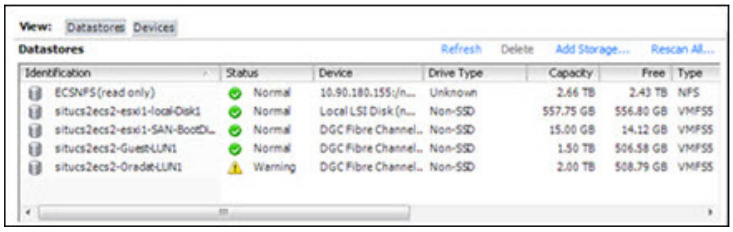

- Oradata-LUN1
- Guest-LUN1
- SAN-BootDisk
- Local-Disk1 (Optional)
- NFS mount to NDS product

Important: ESXi hosts 2 and 4 will NOT have the Oradat-LUN1 datastore.

Modify and Create Port Groups (vSwitch)

Two port groups should be configured on each ESXI host. Each host should include the vSwitch0 and vSwitch1 standard switches. You will edit vSwitch0 and rename the VM network to reflect your subnet. Then create a second vSwitch for the Private Network.

Important: Depending upon you network configuration, you may create additional vSwitches to meet your needs.

Notes: Be sure that the following requirements are met.

- Network IP addresses, as defined by your Network Administrator, have been acquired for the Oracle RAC VMs and ECS VMs.
- A PC exists with vSphere 5.5 or later, PowerShell 3.0, and PowerCLI 5.10.793510 or later installed.

  Note: If you are using vSphere 6.0u1 or later, you will not need PowerShell or PowerCLI.
- The network where the Admin Node resides (includes software and scripts needed for ECS node and RAC installations) is accessible to your PC.
- Credentials (login and password) exist for vCenter, ESXi, and Guest operating system.

Complete the following procedure to modify and create port groups.

1. Select the host, click the **Configuration** tab, and then click **Networking**. The Networking window appears.
2. Click **Properties** for vSwitch0. The vSwitch0 Properties window appears.
3. Select **VM network** and click **Edit**. The VM Network Properties window appears.
4. Change the **Network Label** to the desired name and click **OK**.
5. Click **Close**.
6. Create the second network vSwitch for the private network by clicking **Add networking**.
7. Leave the default values as they are.
8. Select **vmnic1** for this vSwitch and click **Next**.

9 Change the **VM Network** network label to reflect your network.

Note: This example shows a private network created as "RAC_Private".

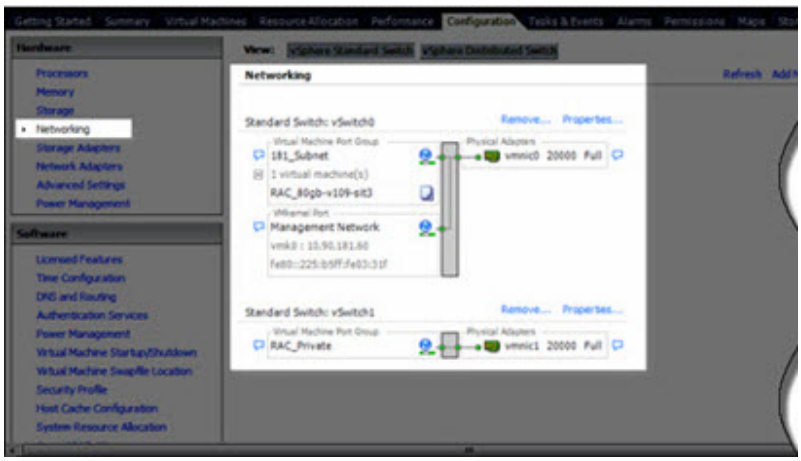

10 Create any additional vSwitches the system needs.

11 Repeat this procedure on each ESXi host.

Important: The Network Label name must be the same for the vSwitch on each ESXi host. It they are not the same, the install can fail.

Creating the RAC Deployment Template

Important: If you are using ESXi 6.0u1 or later, skip this procedure and go to ***RAC Install Using ESXi 6.0u1 or Later*** (on page 143).

Complete the following procedure to create the RAC OVA deployment template.

Note: You should have downloaded the Oracle RAC OVF file from your customer-specific forum on Cisco's File Exchange Server and save it to a local directory. The OVA file is called RAC_16CPU_64GBRAM_45GBHDD_RHEL6U5.ovf.

1 Via a web browser, login to vCenter.

2 Select **ESXi 1 Host** in the datacenter and choose **File > Deploy OVF Template**. The Deploy OVF Template screen is displayed.

3 Click **Browse** and browse to the Oracle RAC OVF file.

Examples:

- This is an example of an OVF file:

  `RAC_16CPU_64GBRAM_45GBHDD_RHEL6U5.ovf`

  Note: Be sure to use the appropriate OVF file for the installation you are performing.

- This is an example path to the OVF file:

  `C:\RAC_install\RAC_16CPU_64GBRAM_45GBHDD_RHEL6U5\RAC_16CPU_64GBRAM_45GBHDD_RHEL6U5.ovf`

Note: Be sure to use your local path to the OVF file.

4 Click **Open**.

5 Click **Next**.

6 Click **Next** again.

7 Name the template and add a suffix to the name to make it unique; then, select the datacenter where the template will be built. Then, click **Next**.

Note: If you will be installing multiple Oracle RAC systems in the same vCenter, add a suffix to the end of the name to make it unique.

Example: This is an example name:

`RAC_16CPU_64GBRAM_45GBHDD_RHEL6U5`

8 Select the **Guest_LUN1** and click **Next** .

9 Select **Thick Provision Lazy Zeroed** and click **Next**.

10 Select the VLAN and click **Next**.

11 Click **Finish**.

12 When the deployment completes, click **Close**.

13 Right-click the VM that was created and choose **Template > Convert to Template**.

14 From the vCenter toolbar, choose **Inventory > VMs and Templates** and verify that the RAC_16CPU_64GBRAM_45GBHDD_RHEL6U5 template is present.

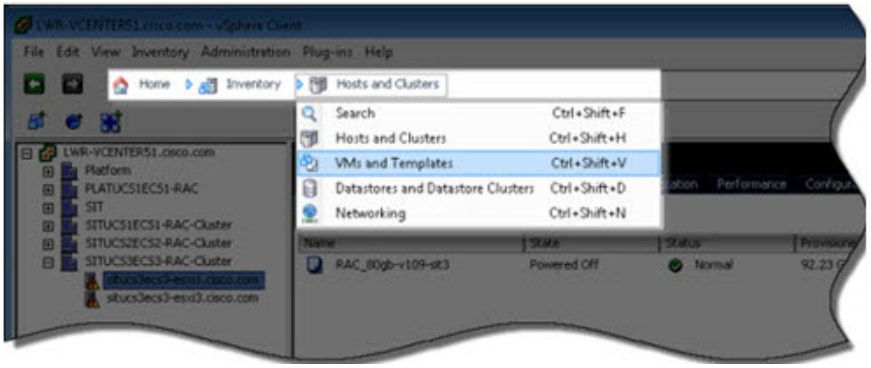

Note: If you will be installing more than one Oracle RAC system on the same vCenter, you may change the template name to specifically identify the system. To do so, right-click the template and select **Rename**. Then, add a suffix to the end of the name.

Example: This is an example name: RAC_16CPU_64GBRAM_45GBHDD_RHEL6U5

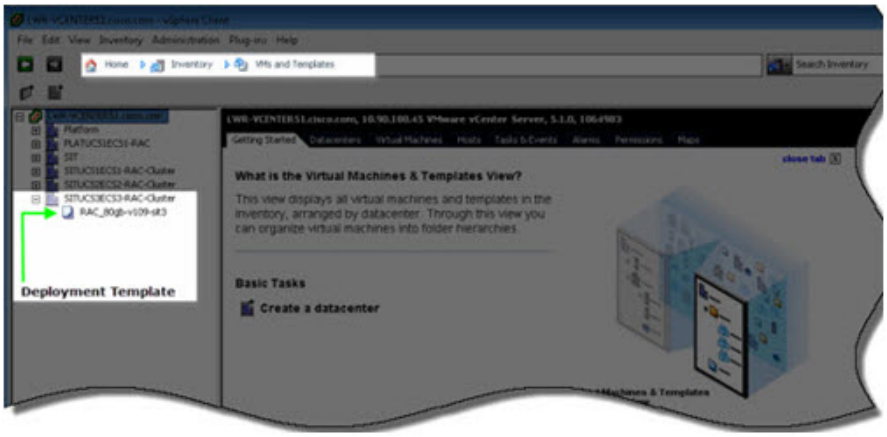

Modifying the Configuration Files

Important: If you are using ESXi 6.0u1 or later, skip this procedure and go to ***RAC Install Using ESXi 6.0u1 or Later*** (on page 143).

1 From the RAC_Install directory on the installation PC, copy the **sample_rac.cfg** file to a new file name that reflects the Oracle RAC system being built.

Example: Copy **sample_rac.cfg** to **rac-site.cfg**.

2 Open the **rac-site.cfg** file in Wordpad.

3 Edit the **rac-site.cfg** file and modify the following entries:

Important: Do NOT change any line that reads "leave empty". These entries will be populated later when the Prepare-OracleRAC_VM.ps1 script is run.

- **HOST_GROUP_NAME** — This name will be used to create the Oracle RAC VMs.

  Example: RAC

- **Database IP Addresses** — Select the IP addresses from the Database network.
  - **PRIMARY_IP** — IP address for access to the VM.
  - **SECONDARY_IP** — IP address for access to the VM.
  - **SCAN_IP_1** — IP address used by services to access the database.
  - **SCAN_IP_2** — IP address used by services to access the database.
  - **VIP_IP_1** — Virtual IP for access.
  - **VIP_IP_2** — Virtual IP for access.
- **RAC Private IP Addresses** — Select the IP addresses from the RAC Private network
  - **MGMT_IP_1** — High Availability management IP address used to communicate between the RAC hosts.
  - **MGMT_IP_2** — High Availability management IP address used to communicate between the RAC hosts.
  - **ETH0_GW** — Gateway IP address.
  - **ETH0_MASK** — Netmask; change this only if it does not meet the needs of the network.
  - **ETH1_MASK** — Netmask; change this only if it does not meet the needs of the network.
- **DATA_DISK_SIZE** = 512 — For ORADATA.

  Note: The DATA_DISK_SIZE should be smaller than the ORADATA LUN.
- **BACK_DISK_SIZE** = 1024 — For ORABACK.

  Note: The BACK_DISK_SIZE should be smaller than the ORABACK LUN.
- **ORACLESID** — This is the default. Do not change.

- **#ORACLE_MEMORY_MODEL —** Do not un-remark unless you want to set the model size shown in the previous chart.
- **AUTO_SIZE_ORACLE_MEM —** This will automatically size the memory. The default is Yes. Leave this field set to **Yes**.
- **NTP_PRIMARY —** Enter the primary NTP IP address.
- **NTP_BACKUP —** Optional for the backup NTP.
- **NTP_ENABLE —** Default is YES. Change to NO if you will NOT use NTP services.
- **ENABLE_RMAN —** Default. Leave this set to YES.

4 Review the changes that you made and save the file.

5 Copy rac-site.cfg to **rac.cfg**.

Note: When entering the file name (rac.cfg) in WordPad, you MUST enclose it in quotes ("). This prevents WordPad from attaching a suffix to the file name (for example, rac.cfg.doc, or rac.cfg.txt).

Installing the Oracle RAC

Important:

- If you are using ESXi 6.0u1 or later, skip this procedure and go to ***RAC Install Using ESXi 6.0u1 or Later*** (on page 143).
- In the event that vCenter is used for multiple datastores, you must ensure that your naming schema is unique for each datastore and its hosts. If you use the same naming schema, your latest install will overwrite the existing server and may even fail.

Complete the following procedure if you are using ESXi 5.5 or 6.0 to install the Oracle RAC for your NextX system.

1 Right-click the PowerShell application and select **Run as Administrator**.

2 Change to the directory location where the Prepare-OracleRAC_VM_RHEL6_v_2_07A.ps1 file exists.

Example: cd c:\RAC_install

3 Execute the **Prepare-OracleRAC_VM_RHEL6_v_2_07A.ps1** script as shown in the following example.

Example: .\Prepare-OracleRAC_VM_RHEL6_v_2_07A.ps1

4 When prompted, enter your vCenter login and password credentials. You are prompted for the Oracle RAC you are installing.

Note: Ignore any warnings that may appear.

Important: If you are not executing the script from the vCenter PC, you are prompted for the vCenter IP address. Enter the **IP address** and press **Enter**.

5 When prompted for the template you want to use for the Oracle RAC, enter the item number associated with your template (e.g. RAC_16CPU_64GBRAM_45GBHDD_RHEL6U5).

```
1. PLATUCS1ECS1_RAC1_16CPU_64GBRAM_80GBHDD-v109
2. PLATUCS1ECS1-RAC2_16CPU_64GBRAM_80GBHDD-v109
3. RAC_16CPU_64GBRAM_45GBHDD_RHEL6U5
```

6 When prompted for the datacenter, enter the item number associated with the datacenter you are deploying and press **Enter**.

7 When prompted for the first esxhost, enter the item number associated with esxi-1 and press **Enter**.

Note: The first esxhost is the Blade-1 host.

Example: situcs2ecs2-esxi1

8 When prompted for the second esxhost, enter the item number associated with esxi-3 and press **Enter**.

Note: The second esxhost is the Blade-3 host.

Example: situcs2ecs2-esxi3

9 When prompted for the datastore where the Guest operating system disk will be deployed, select the appropriate item number (for example, the number representing [datastore] **Guest-LUN**) and press **Enter**.

Important: Two File Manager windows will appear, one at a time. You will be prompted to select specific ISO files. Do NOT cancel these windows; otherwise, you will need to close out of your PowerShell session and restart the Prepare-OracleRAC_VM_RHEL6_v_2_07A.ps1 script.

10 A file manager window will open. Select the **RAC_12c_rpms_RHEL6.x_EE ISO** and click **Open**.

Important: Do not close this window. You must select the rpms ISO and click **Open**.

11 A file manager window will open. Select the following ISO and click **Open**.

- **RAC_addon_scripts_RHEL6.x_EE_v2.0.7**

Important: Do not close this window. You must select the rmps ISO

Note: A message stating the files were found should appear. Enter **N** if files exist.

12 When prompted for the port group to use for the Management Network (eth0), enter the appropriate item number and press **Enter**.

Note: For this sample install, our network is 181_Subnet.

13 When prompted for the port group to use for the private network (eth1), enter the appropriate item number and press **Enter**.

14 When prompted for the datastore for the voting disk, select the item number for the "Oradat" datastore and press **Enter**.

Example: In our installation example, the Oradat datastore is "situcs2ecs2-Oradat-LUN1".

```
1. ECSNFS
2. situcs2ecs2-Guest-LUN1
3. situcs2ecs2-Oradat-LUN1
------------------------
Please enter the number of the datastore for the voting disk: 3
```

15 When prompted for the datastore for the Oracle database, select the item number for the "Oradat" datastore and press **Enter**.

16 When prompted for the datastore for the Oracle backup, select the item number for the "Oradat" datastore and press **Enter**.

17 When prompted for the ESXi root password, enter the password and click **OK**.

Note: Enter the ESXi password for this system.

18 When prompted for the VM Guest root password, enter the password and click **OK**. The installation begins.

Important:

- For the Guest root password, enter the default root password. You will change the password later in this document.
- Once the installation begins, do NOT click in the PowerShell window, because it will pause the script. If you inadvertently pause the script, click in the PowerShell window and press **Enter** to continue.

19 This script will take at least 1 to 2 hours to complete.

Note: If, for some reason, the script fails, you MUST check the rac.cfg file to ensure that the "leave blank" fields were not populated. If they were, you must copy the sample.cfg file (created earlier) back to rac.cfg before re-executing the Prepare script.

20 Periodically monitor the progress of the installation from the PowerShell progress bar located at the top of the window. When this completes, post install scripts are executed.

Note: The script will display a message indicating when it has completed.

21 In the PowerShell window, type **exit** and press **Enter**. The PowerShell window closes

RAC Installation Verification Procedures

1. From vCenter, verify that the ESXi 1 and 3 hosts include a VM with the name specified in the rac.cfg file. The ESXi 1 VM should have an "a" appended at the end of the name; ESXi 3 host should have a "b" appended at the end of the name.
   Example: sitecs2-RAC2a and sitecs2-RAC2b
   Important: The names are unique to this installation.

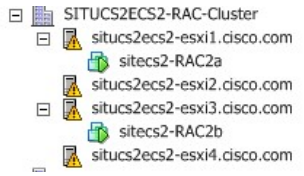

2. In vCenter, click the ESXi 1 VM and then click the **Console** tab.
3. Click in the screen and press **Enter** to display the login prompt.

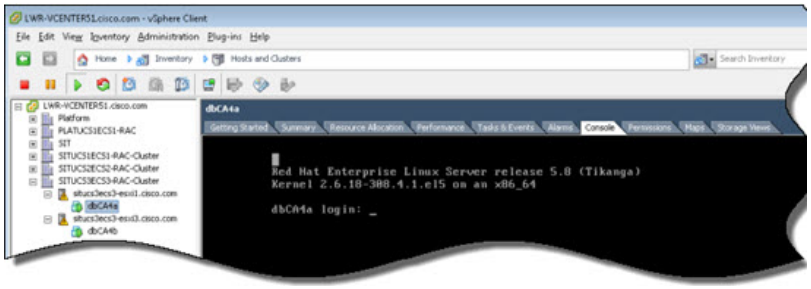

4. Log in with the Guest **root** credentials.
5. At the prompt, enter the following command to redirect the output of your network interface configuration to a file.

```
ifconfig -a > /tmp/ifconfig.out
```

6. Review the **/tmp/ifconfig.out** file. You should have the following network interfaces configured.
   Note: The IP addresses will be unique to your installation.

```
eth0      Link encap:Ethernet  HWaddr 00:50:56:99:41:FB
          inet addr:203.0.113.243  Bcast:203.0.113.255  Mask:255.255.255.0
eth0:1    Link encap:Ethernet  HWaddr 00:50:56:99:41:FB
          inet addr:203.0.113.247  Bcast:203.0.113.255  Mask:255.255.255.0
eth0:3    Link encap:Ethernet  HWaddr 00:50:56:99:41:FB
          inet addr:203.0.113.246  Bcast:203.0.113.255  Mask:255.255.255.0
eth1      Link encap:Ethernet  HWaddr 00:50:56:99:2B:A4
          inet addr:198.51.100.1  Bcast:198.51.100.255  Mask:255.255.255.0
eth1:1    Link encap:Ethernet  HWaddr 00:50:56:99:2B:A4
          inet addr:169.254.116.48  Bcast:169.254.255.255  Mask:255.255.0.0
          inet addr:127.0.0.1  Mask:255.0.0.0
```

7 Repeat Steps 2 through 6 for the VM on ESXi host 3. The output of the ifconfig file should have the following network interfaces configured.

Note: The IP addresses will be unique to your installation.

```
eth0      Link encap:Ethernet  HWaddr 00:50:56:99:01:7A
          inet addr:203.0.113.244  Bcast:202.0.113.255  Mask:255.255.255.0
eth0:1    Link encap:Ethernet  HWaddr 00:50:56:99:01:7A
          inet addr:203.0.113.245  Bcast:203.0.113.255  Mask:255.255.255.0
eth0:2    Link encap:Ethernet  HWaddr 00:50:56:99:01:7A
          inet addr:203.0.113.248  Bcast:203.0.113.255  Mask:255.255.255.0
eth1      Link encap:Ethernet  HWaddr 00:50:56:99:56:55
          inet addr:198.51.100.2  Bcast:198.51.100.255  Mask:255.255.255.0
eth1:1    Link encap:Ethernet  HWaddr 00:50:56:99:56:55
          inet addr:169.254.80.73  Bcast:169.254.255.255  Mask:255.255.0.0
          inet addr:127.0.0.1  Mask:255.0.0.0
```

8 Verify that each server can ping each IP address of the other server.

9 Return to the VM on ESXi host 1 and review each installation log in **/var/log/nds/ora_cluster**.

10 Review the following root log files in the following directory:

/opt/oracle/installed/oracle_cluster-12.1.0.2-0/install

11 Next, review the database-build installation logs in **/var/log/nds/cabhe**.

12 Review the logs in the **c:\RAC_install** directory. The logs will have filenames similar to the following:

- **RAC_install_<data_time>**
- **RAC_Install_GuestOS_Output-<date_time>**

13 On each VM for ESXi host 1 and ESXi host 3, enter the following command to verify that the Oracle processes are running.

```
ps -ef | grep -i ora
```

14 Execute the following commands on the ESXi host 1 VM to verify the system configuration:

```
source /opt/oracle/CRS.env
crsctl status resource -t | more
```

Example Output: There will be a lot of output. Verify all TARGET and STATE values match (ONLINE / ONLINE). Also verify the following:

```
ora.asm
```

ONLINE ONLINE dbca3a Started,STABLE

ONLINE ONLINE dbca3b Started,STABLE

ora.oraback.oraback_vol.acfs

ONLINE ONLINE platucs1ecs1-rac1a mounted on /oraback,STABLE

ONLINE ONLINE platucs1ecs1-rac1b mounted on /oraback,STABLE

ora.cabhe.db

1 ONLINE ONLINE dbca3a Open,STABLE2 ONLINE ONLINE dbca3b Open,STABLE

Note: Ensure the TARGET and STATE columns match for each NAME.

15 Enter the following command to verify that the RAC is functioning properly.

```
crsctl status resource -t
```

16 The Oracle installation is complete. Go to the next procedure in this chapter.

Add Another Interface to the RAC Nodes (Optional)

Important: If you are using ESXi 6.0u1 or later, skip this procedure and go to Verifying the Enterprise Manager.

1 From the vSphere Web UI, right-click the **RAC VM** and click **Edit Settings**.

2 Click **Add...**.

3 Select the **Ethernet Adapter** and click **Next**.

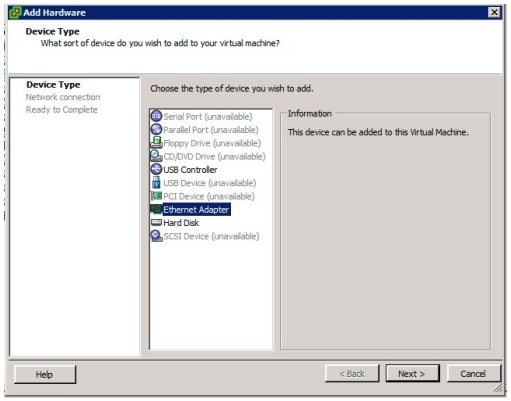

4 Leave the **Adapter Type** at the default value (E1000).

5 Select the correct network under **Network Connection**.

6 Select **Connect at power on**.

7 Click **Next**.

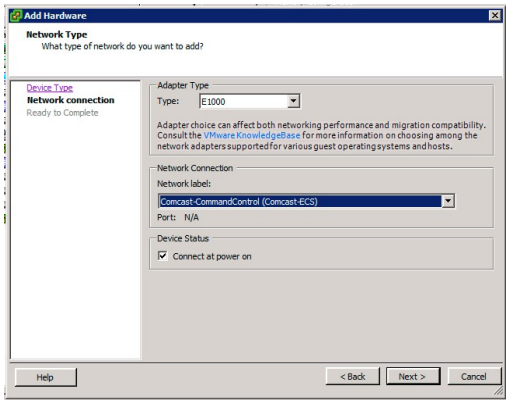

8 Click **OK** on the confirmation screen.

9 Click **OK** on the Virtual Machine Properties screen. A message similar to the following appears on the VM's console:

```
PCI: Enabling device 0000:02:03.0 (0100 -> 0103)
```

Note: This indicates that the VM OS recognizes the new network adapter.

10 Enter the following command to see more detailed information.

```
[root@nextxecs4rac4a ~]# dmesg | tail
```

Note: The last six lines show the more relevant information.

```
[Oracle OKS]     Node 1 (IP 0xa9fe8dde) Node 2 (IP 0xa9fe2a79)
[Oracle OKS]  Node count 2, Local node number 1
ADVMK-00013:  Cluster reconfiguration started.
ADVMK-00014:  Cluster reconfiguration completed.
[0]:  VMCI:  Updating contect from (ID=0xffffffff) to (ID=0xb73ba6bb) on event (type=0).
PCI:  Enabling device 0000:02:03.0 (0100 -> 0103)
ACPI:  PCI Interrupt 0000:02:03.0[A] -> GSI 17 (level, low) -> IRQ 51
PCI: Setting latency timer of device 0000:02:03.0 to 64
e1000: 0000:02:03.0: e1000_probe: (PCI:66MHz:32-bit) 00:50:56:bs:dd:9e
e1000: eth2: e1000_probe: Intel(R) PRO/1000 Network Connection
```

Configure the New Network on the RAC Node

Complete the following steps to configure the new network.

1 As **root** user on the *primary* RAC node, enter the following commands to create the **/etc/sysconfig/network-scripts/ifcfg-eth2** file.

```
[root@nextxecs4rac4a ~]# cd /etc/sysconfig/network-scripts
```

2 Enter the following command to copy the ifcfg-eth0 file to **ifcfg-eth2**.

```
[root@nextxecs4rac4a ~]# cp -i ifcfg-eth0 ifcfg-eth2
```

3 Enter the following command to open the ifcfg-eth2 file in a text editor.

```
[root@nextxecs4rac4a ~]# vi ifcfg-eth2
```

4 Modify the DEVICE, IPADDR, NETMASK, and HWADDR fields to reflect network values for the new network interface.

Note: Setting the HWADDR field is a way of avoiding issues where eth1 and eth2 swap inside RedHat during bootup.

Example:

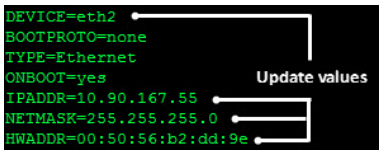

5 Save and close the file.

6 Open the **/etc/sysconfig/network** file in a text editor and modify the GATEWAY field, as needed.

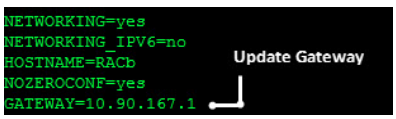

7 Save and close the file.

8 Reboot the RAC node.

9 Repeat these steps [beginning with ***Add Another Interface to the RAC Nodes (Optional)*** (on page 24)] for the *secondary* RAC node.

Increasing the Oracle Database Process Limit

Complete this procedure to increase the number of database processes from 1000 to 2000.

1. Log into the Oracle RAC Server node as **root** user.
2. Type the following commands and press **Enter** to change to the oracle user and to source the CABHE environment:

   ```
   [root@nextxecs4rac4a ~]# su - oracle
   [oracle@nextxecs4rac4a ~]$ . /opt/oracle/CABHE.env
   ```

   Note: There is a space between the "." and "`/opt`".
3. Log into the database as **sysdba** user.

   ```
   [oracle@nextxecs4rac4a ~]$ sqlplus / as sysdba
   ```
4. Type the following sql statement and press **Enter** to create the pfile from the database.

   Note: Replace <Database_Name> with the actual database name for your system.

   Command Syntax:

   ```
   create pfile= '/oraback/<Database_Name>/
   pfile-change-process.txt ' from spfile;
   ```

   Example:

   ```
   SQL> create pfile= '/oraback/CABHE/pfile-change-process.txt
   ' from spfile;
   ```
5. Type the following sql statement and press **Enter** to obtain the spfile location.

   Command:

   ```
   select VALUE FROM v$parameter WHERE NAME='spfile';
   ```

   Example Output:

   ```
   SQL> +ORADATA/cabhe/spfilecabhe.ora
   ```
6. Enter the following command to set the processes to **2000** and to define the scope.

   ```
   SQL> alter system set processes=2000 scope=spfile;
   ```
7. Type **exit** and press **Enter** to log out of the database.
8. Type the following commands and press **Enter** to shutdown the database.

   ```
   [oracle@nextxecs4rac4a ~]$ srvctl stop database -d CABHE
   [oracle@nextxecs4rac4a ~]$ srvctl start database -d CABHE
   [oracle@nextxecs4rac4a ~]$ srvctl status database -d CABHE
   ```

9 Type the following commands and press **Enter** after each to check the number of processes and sessions.

```
[oracle@nextxecs4rac4a ~]$ sqlplus / as sysdba
SQL> select name, value from v$parameter where name in
('processes', 'sessions', 'transactions');
```

Note: This select statement is one contiguous line. Do not press Enter until you have typed the entire command.

Example output:

NAME	VALUE
processes	2000
sessions	3024
transactions	3326

Enable Oracle Database Backups

In this procedure, you will check the Oracle database archive mode and modify the following if necessary:

- Oracle DB Archive
- RMAN for database backups
- Automated database backups in crontab

Verifying that the Oracle Database is in Archive Mode

Complete these steps to verify that the Oracle database is in Archive mode.

Note: If you have completed all procedures and steps to this point, you should still logged into the Oracle database.

1 Type the following command and press **Enter** to check if Oracle is in Archive mode:

```
SQL> archive log list;
```

Sample output: The response should be as follows when the database log mode set to Archive mode:

```
Database log mode              Archive Mode
Automatic archival             Enabled
Archive destination            USE_DB_RECOVERY_FILE_DEST
Oldest online log sequence     137
```

Note: If the database is NOT in Archive mode, a response similar to the following appears:

```
Database log mode              No Archive Mode
Automatic archival             Disabled
Archive destination            USE_DB_RECOVERY_FILE_DEST
Oldest online log sequence     137
Current log sequence           140
Next log sequence to archive   140
```

2 Is the database in Archive mode?

- If **yes**, type `exit` and press **Enter** to exit the database. You are finished with this procedure. Go to *Verifying Oracle RAC RMAN Presence* (on page 31).
- If **no**, go to *Changing an Oracle Database to Archive Mode - RAC Installation* (on page 29).

Changing an Oracle Database to Archive Mode - RAC Installation

Follow this procedure only if the database is not in Archive mode.

Important: If you are working on a production database, get permission from the database administrator (DBA) before running the following procedure.

1. At the SQL prompt, type `exit` and press **Enter** to exit the database.
2. Type the following command and press **Enter** to verify that the **/oraback** file system is mounted.

   ```
   [oracle@nextxecs4rac4a ~]$ df -h
   ```

3. Verify that the following directory structure exists.

   Note: Substitute the database name for <Database_Name> in the command syntax below.

   ```
   /oraback/<Database_Name>
   ```

   Examples:

   ```
   ls -l /oraback/CABHE
   drwxr-xr-x 2 oracle dba 4096 Jun 10 17:56 archive
   drwxr-xr-x 2 oracle dba 4096 Jun 10 17:56 export
   drwxr-xr-x 2 oracle dba 4096 Jun 10 17:56 rman
   ```

   Note: The CABHE and archive directories must be owned by oracle, group dba.

4. Enter the following command to source the **CRS** environment for the oracle server.

   ```
   [oracle@nextxecs4rac4a ~]$ source /opt/oracle/CRS.env
   ```

5. Type the following command and press **Enter** to shut down the database instances.

   ```
   [oracle@nextxecs4rac4a ~]$ srvctl stop database -d CABHE
   ```

6. Type the following command and press **Enter** to verify that the instances are down.

   ```
   [oracle@nextxecs4rac4a ~]$ srvctl status database -d CABHE
   ```

   Note: You should see output similar to the following:

   ```
   Instance CABHE01 is not running on node <RacHostA>
   Instance CABHE02 is not running on node <RacHostB>
   ```

7. Type the following command and press **Enter** to source the Oracle database environment:

   Command Syntax:

   ```
   source /opt/oracle/<SID>.env
   ```

   Example

   ```
   [oracle@nextxecs4rac4a ~]$ source /opt/oracle/CABHE.env
   ```

8 Type the following command and press **Enter** to access the database:

```
[oracle@nextxecs4rac4a ~]$ sqlplus / as sysdba
```

9 Type the following command and press **Enter** to start a single instance of the database:

```
SQL> startup mount;
```

10 Type the following command and press **Enter** to change the database to Archive mode:

```
SQL> alter database archivelog;
```

11 Type the following command and press **Enter** to change the configuration of the database to point to the directory location for all archives:

Notes:

- This procedure establishes the archive location. Take into consideration the amount of space you might need in the future to hold all archive logs.
- This command may fail the first time you execute it. If it does, execute the command again.

```
SQL> alter system set
log_archive_dest_1='LOCATION=/oraback/CABHE/archive/';
```

12 Start the Oracle server.

```
SQL> alter database open;
```

13 Type `exit` and press **Enter** to exit the database.

14 Type the following command and press **Enter** to restart the database for all instances:

```
[oracle@nextxecs4rac4a ~]$ srvctl start database -d CABHE
```

15 Type the following command and press **Enter** to check the database status. An instance should be running on both servers.

```
[oracle@nextxecs4rac4a ~]$ srvctl status database -d CABHE
```

Note: You should see output similar to the following:

```
Instance CABHE01 is running on node <RacHostA>
Instance CABHE02 is running on node <RacHostB>
```

16 Type the following command and press **Enter** to verify the database is in Achive mode:

```
[oracle@nextxecs4rac4a ~]$ sqlpus / as sysdba
SQL> archive log list;
```

17 Does the output show that the database log mode is set to **Archive** mode?

- If **yes**, go to the next step.
- If **no**, repeat this procedure. If the database still does not set the database log mode to Archive mode, contact Cisco for assistance.

18 Type `exit` and press **Enter** to exit the database.

19 Type `exit` and press **Enter** to log out of the oracle user.

Verifying Oracle RAC RMAN Presence

The RMAN rpm gets installed as part of the Oracle RAC installation.

Note: This procedure must be executed on *both* the RAC nodes.

1 Login to the *primary* RAC node as **root** user.

2 Type the following command and press **Enter** to verify that the RMAN rpm exists.

```
[root@nextxecs4rac4a ~]# rpm -qa | grep rman
```

Example Output:

```
nds_ora_rman-1.0.4-0_el6_11g.x86_64
```

3 Change the permission for **/var/log/nds** directory to **755**.

```
[root@nextxecs4rac4a ~]# chmod 755 /var/log/nds
[root@nextxecs4rac4a ~]# ls -ltrd /var/log/nds
drwxr-xr-x. 18 ndsuser ndslog 4096 Nov 17 17:00 nds
```

4 Repeat this procedure on the *secondary* RAC node.

Directory Tree for RMAN

The RMAN directory structure should resemble the example under the **/opt/nds** directory.

1 Type the following commands and press **Enter** to verify that a link to rac_ora_rman exists.

```
[root@nextxecs4rac4a ~]# cd /opt/nds
[root@nextxecs4rac4a nds]#  ls -l  ora_rman
```

Expected output:

```
ora_rman -> /opt/nds/installed/ora_rman-1.0.4-0
```

2 Go to the ora_rman directory and list the contents of the directory.

```
[root@nextxecs4rac4a ora_rman]# cd ora_rman
[root@nextxecs4rac4a ora_rman]#  ls -l *
```

Notes:

- The directory structure under the ling rac_ora_rman should look like the list below.

  |-- bin

  |-- docs

  |-- sample

  |-- etc

  |-- licenses

  |-- log

|-- sql

|-- utils

Note: The files under /opt/nds/ora_rman/utils are required for running RMAN backup and restore.

Modifying the RMAN Configuration File for Backups

Before running RMAN backups, you must copy, rename, and edit the CABHE.rman.config file.

Important: You must copy and edit the CABHE.rman.config file on both Oracle RAC nodes.

1. If necessary, login to the *primary* RAC node as the **root** user.
2. Type the following command and press **Enter** to copy the CABHE.rman.config file to the **/opt/nds/rac_ora_rman/etc** directory.

   ```
   [root@nextxecs4rac4a ~]# cd /opt/nds/ora_rman/docs/sample
   [root@nextxecs4rac4a sample]# cp -p CABHE.rman.config
   ../../etc
   ```

3. Type the following command and press **Enter** to change to the required directory.

   ```
   [root@nextxecs4rac4a sample]# cd ../../etc
   ```

4. Edit the **CABHE.rman.config** file in a text editor.

   ```
   [root@nextxecs4rac4a etc]# vi CABHE.rman.config
   ```

5. Update the following variables to the values shown below.

   Note: Keep the double quotes.

   ```
   SID_NAME="CABHE01"
   BACKUP_VERSIONS="3"
   TRANSFER_TO_STANDBY="N"
   DIRECTORY_LOC_BACKUP="/oraback/CABHE/rman"
   DIRECTORY_LOC_LOG="/opt/nds/rac_ora_rman/log"
   ```

6. Save and close the file.

   Note: If you want to backup to a remote (NFS) storage device, create an NFS mount entry in the /etc/fstab file. Change the SECOND_DIRECTORY="Y" and change the path of the DIRECTORY_LOC_BACKUP2 variable to the mount point used for the remote (NFS) storage device.

7. Repeat Steps 1 through 6 on the *secondary* RAC node and set the variables in the **CABHE.rman.config** file to the following values.

   ```
   SID_NAME="CABHE02"
   BACKUP_VERSIONS="3"
   TRANSFER_TO_STANDBY="N"
   SECOND_DIRECTORY="Y"
   DIRECTORY_LOC_BACKUP2="[/nfs/mount/path]"
   ```

Note: Replace [/nfs/mount/path] with the actual mount path to your NFS device.

8 Save and close the **CABHE.rman.config** file.

Manually Execute Database Backups

Complete the following steps to verify that the a database backups run successfully.

Note: For more detailed information about database backups, refer to the ***Explorer Controller Suite 3.0 Backup and Restore User Guide***.

1 Enter the following command on either RAC node to execute a full database backup. When the backup completes, a **Backup has completed successfully** message appears.

```
[root@nextxecs4rac4a ~]#
/opt/nds/ora_rman/utils/ora_rman_wrapper.sh -f CABHE
```

```
Hostname:  nextxecs4rac4a

Service Name:  nds_oracle_cabhe
No Cluster Service on this server
No Cluster Service on this server

Calling Full Backup
OWNER = nextxecs4rac4a
HOSTNAME = nextxecs4rac4a
STATUS = started

Starting Full Backup
Fri Jul 21 13:01:20 Full Backup has started on <HOSTNAME>
================================================================================
You can monitor progress on Log file /opt/nds/ora_rman/log/CABHE_backup_Full_21-
Jul-2017.log
================================================================================
Fri Jul 21 13:02:15 Backup has completed successfully
```

2 Enter the following command to test a differential backup. When the backup completes, a **Differential Backup has completed successfully** message appears.

```
[root@nextxecs4rac4a ~]#
/opt/nds/ora_rman/utils/ora_rman_wrapper.sh -d CABHE
```

```
Hostname:  nextxecs4rac4a

Service Name:  nds_oracle_cabhe
No Cluster Service on this server
No Cluster Service on this server

Calling Diff Backup

Starting Diff Backup
Fri Jul 21 13:10:17 Differential Backup has started on <HOSTNAME>
================================================================================
You can monitor progress on Log file /opt/nds/ora_rman/log/CABHE_backup_Diff_21-J
ul-2017.log
================================================================================
Fri Jul 21 13:11:11 Differential Backup has completed successfully
```

Adding a cron Job for Automated Database Backups

Important: You must execute this procedure on both RAC nodes.

1 As **root** user on the *primary* RAC node, type the following command and press **Enter** to edit the root crontab file.

```
[root@ecs4rac4a ~]# crontab -e
```

2 Add the following two cron jobs to the file.

```
# Perform full database backup at 2 AM every Sunday
0 2 * * 0 /opt/nds/ora_rman/utils/ora_rman_wrapper.sh -f CABHE
> /opt/nds/ora_rman/log/full_database_backup_cron.out 2>&1
# Perform differential database backup at 2AM every other day
of week
0 2 * * 1-6 /opt/nds/ora_rman/utils/ora_rman_wrapper.sh -d
CABHE > /opt/nds/ora_rman/log/diff_database_backup_cron.out
2>&1
```

3 Save and close the file.

4 Have you completed this procedure on both RAC nodes?

- If **yes**, go to the next procedure in this document.
- If **no**, repeat Steps 1 through 3 on the *secondary* RAC node.

  Note: Although Step 1 states the primary RAC node, you will execute this step on the *secondary* RAC node.

Change the root and oracle User Passwords

In this procedure, you will change the default root and oracle passwords to a site-specific, secure password. Complete the following steps to change these passwords.

Complete the following steps to change the root and oracle passwords on your RAC nodes.

Important: The passwords MUST conform to these restrictions:

- A valid password should be a mix of upper-case and lower-case letters, digits, and special characters.
- The password must contain at least one special character (**example:** ! $ @).
- Your password must be 10 characters long, with characters from at least 3 of these above-mentioned classes.

1. Log into the *primary* Oracle database VM using the default **root** password.
2. Type the following command and press **Enter**.

   ```
   [root@nextxecs4rac4a etc]# passwd root
   ```

   Result: The **Changing password for user root** message displays with the password requirements.
3. Type a password that conforms to the password restrictions and press **Enter**. If the password meets restrictions, you are prompted to retype the new password.

   Note: If the password does not meet restrictions, you are prompted to enter a new password. Enter a password that meets all of the restrictions.
4. Retype the same password and press **Enter**. If both entries match, the password is changed.

   Notes:

   - If the password change is successful, the **passwd: all authentication tokens updated successfully** message displays.
   - If the passwords do not match, the **Sorry, passwords do not match** displays and you are prompted to enter a new password.
5. Was the password successfully changed?

   - If **yes**, you are returned to the root prompt. Continue with the next step to change the oracle user password.
   - If **no**, repeat steps 3 and 4.

6 Type the following command and press **Enter** to change the oracle user password.

```
# password oracle
```

7 Enter a new password and press **Enter**.

8 Retype the new password and press **Enter**.

9 Was the password successfully changed?

- If **yes**, you are returned to the root prompt.
- If **no**, repeat steps 7 and 8.

10 Repeat Steps 1 through 9 on the *secondary* RAC node.

3

Create ECS Database Users

This chapter includes the procedure to download the script to create Oracle database users for your ECS. It also includes the steps to create the database users required for your ECS system.

In This Chapter

Downloading the ECS Database User Script

The create_ora_project_ASM.sh script is available on the Admin Node as it is package in the cisco-vcs-deployment.[VERSION]zip file. Typically the script will reside in the following directory on the admin node: /opt/cisco/software/admin_node/cisco-vcs-deployment-[version]/scripts/.

Complete the following steps to copy the create_ora_project_ASM.sh script to one of your RAC nodes.

1 As **root** user on one of the RAC nodes, enter the following command to change to the **/home/oracle** directory.

```
[root@nextxecs4rac4a etc]# cd /home/oracle
```

2 Using SCP, copy the **create_ora_project_ASM.sh** file from the Admin Node to the /home/oracle directory.

Note: Substitute the IP address for one of your RAC nodes in the command below.

Command Syntax:

```
scp admin@[Admin_Node_IP]:/opt/cisco/software/admin_node/
cisco-vcs-deployment-1.0.6/scripts/create_ora_project_ASM.sh .
```

Example:

```
[root@nextxecs4rac4a ~]# scp
admin@[10.90.44.70]:/opt/cisco/software/admin_node/
cisco-vcs-deployment-1.0.6/scripts/create_ora_project_ASM.sh .
```

3 Enter the following command to verify that the script is present in the /oracle/home directory.

```
[root@nextxecs4rac4a ~]# ls -ltr /home/oracle
```

Example Output:

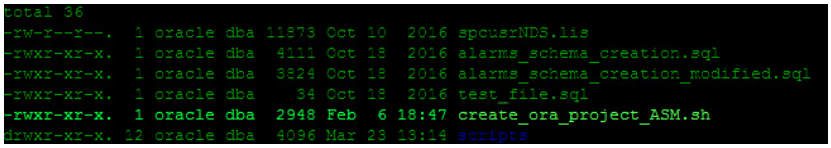

4 Go to the next section to create database user.

Creating Database Users

Complete the following steps to create the database users for your NextX system.

1 As **root** user, change to the **oracle** user.

```
[root@nextxecs4rac4a ~]# su - oracle
```

2 Copy the **create_ora_project_ASM.sh** script to the Oracle home directory.

Note: There is a space between the asterisk (*) and the period.

```
[oracle@nextxecs4rac4a ~]$ cp /home/oracle/create* .
```

3 Enter the following command to setup the oracle environment that is required to run the script.

```
[oracle@nextxecs4rac4a ~]$ CABHE
```

4 Enter the following command to execute the **create_ora_project_ASM.sh** script to create the following three database users.

- VCS Console user
- ECS user
- ALARMS user

Command Syntax:

```
./create_ora_project_ASM.sh <project name>
```

Example:

```
[oracle@nextxecs4rac4a ~]$ ./create_ora_project_ASM.sh vcsprod
```

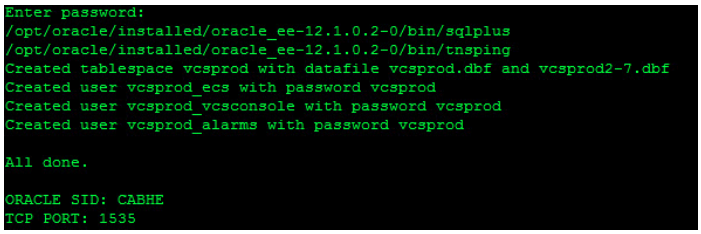

5 When prompted, enter a password. This password will be set for all three database users.

Important: The database user password must be 6 or more characters. This example uses vcsprod as the password. which is the same as the project name.

6 Log back into the Oracle database as **sysdba** and then enter the following command to verify that the database users were successfully created.

Important: Replace VCSPROD with the <project name> you enter in Step 5.

```
[oracle@nextxecs4rac4a ~]$ sqlplus / as sysdba
SQL> select username from all_users where username like
'%VCSPROD%';
```

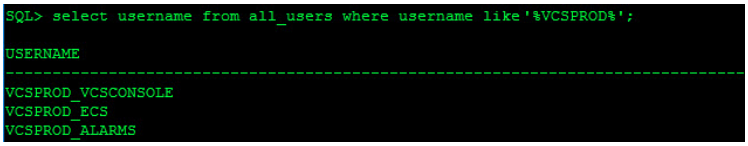

7 Record the database user names as they will be needed in other procedures throughout this document.

4

Install and Configure the ECS System

Important: Complete the procedures in the section for either a new installation or an ECS migration. The procedures must be executed in the order in which they are written.

This chapter provides the procedures to build the entire ECS system using the Linux platform template that was built when the Admin Node was created. You will need the following information to continue.

- Location of the Linux platform template
- The admin user password defined for the Linux platform template
- NextX X.509 Root CA and NextX X.509 Certificates
- The Root CA certificate and the NextX X.509 certificates for each ECS node should have been created when you deployed the Admin Node. If they have not been created, refer to **Chapters 5 and 6** of the ***Admin Node Installation Guide*** to create them now.

In This Chapter

Install the Consul VM

Important: A total of three Consul VMs will be deployed as two of the VMs are used for failover purposes. You will need to repeat all of the procedures in this section three times.

The Consul VM is used as a tool for discovering and configuring services in the ECS infrastructure, such as service discovery and health checks. Complete the procedures in this section to deploy the Consul VMs using the Linux platform template.

Please ensure that your system meets the following requirements before proceeding.

Important: If any of these requirements have not been completed, please do so now.

- Linux platform template built
- Admin Node deployed and running
- NextX X.509 Root Certificate Authority (CA) Certificates are created on the Admin Node
- Three IP addresses are available from your System Administrator
  - Consul 1 IP
  - Consul 2 IP
  - Consul 3 IP

Creating the Consul VM

Important: If you are using vSphere client to deploy virtual machines, you cannot create the Consul VM using a template. Refer to ***Procedures When Using vSphere Client*** (on page 221).

Complete the following procedures to deploy the Consul VM from the vSphere Web UI.

1. From the vCenter Web UI, click **VMs and Templates**.
2. Locate and select the CSCOlxplat template that was built using the ***Admin Node Installation Guide***.

3 Right-click the template and select **Deploy VM from this Template**. The Deploy From Template window opens.

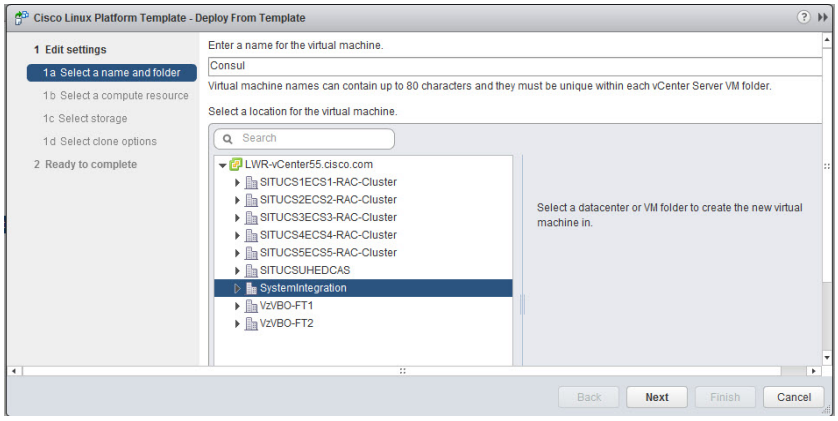

4 In the text box, enter a name for the VM you are creating and then select the datacenter where it will be deployed. Click **Next.**

5 Select the appropriate ESXi host where you want to deploy the VM and click **Next**.

6 From the **Select virtual disk format** dropdown menu, maintain the **Same format as source** default. Then ensure that the appropriate datastore is selected.

7 Click **Next** and then click **Next** again.

8 Review the settings and click **Finish**.

Reconfiguring the Consul VM

Important: If you are using vSphere client to deploy and configure VMs, refer to ***Reconfiguring the Virtual Network Using vSphere Client*** (on page 223).

Complete the following procedure to reconfigure the virtual hardware on the Consul node.

1 From the vSphere Web UI, click the **Home** icon, , and then click **Hosts and Clusters**.

2 Locate and select the Consul VM.

3 Right-click the VM and select **Edit Settings**. The Edit Settings window appears.

4 From the **CPU** text box, change the value to **1**.

5 From the **Memory** text box, change the value to **2 GB**.

6 From the **Hard disk 1** text box, change the value to **32 GB**.

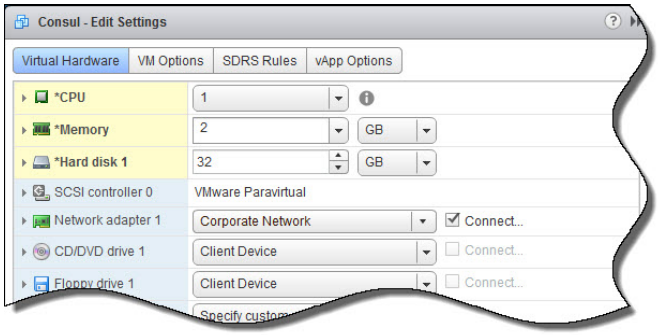

7 Click **OK**. The VM is reconfigured.

8 Monitor the **Recent Tasks** area to confirm that the VM virtual hardware is successfully reconfigured.

Configuring the Consul Network Interface With a Static IP

Complete the following steps to reconfigure the network interface with a static IP address.

Note: Your network administrator should have provided you with a static IP address, default gateway and a network mask value.

1 Select and right-click the Consul VM and select **Power On**.

2 Select and right-click the Consul VM again and select **Open Console**. A VMware console window opens in a new tab.

3 Log into VM with the following credentials.

Username: admin

Password: [password for Linux platform template]

4 Type the following command line utility to configure the network and the DNS settings. The Select Action window appears.

```
[admin@platform ~]$ sudo system-config-network
```

5 Select **Device configuration**.

6 Highlight **eth0** and press **Enter**.

7 Press the **Tab** key until you **Use DHCP i**s highlighted. Then press the **Spacebar** to unselect this option.

8 Tab to each field to enter the following values.

Note: Domain Name Server (DNS) entries are optional.

- **Static IP**
- **Netmask**
- **Default gateway IP**
- **Primary DNS Server**
- **Secondary DNS Server**

9 Verify that **Peer DNS** is selected.

10 Press the **Tab** key until you **Controlled by NetworkManager i**s highlighted. Then press the **Spacebar** to unselect this option.

11 Press **Tab** to highlight **Ok** and press **Enter**. The Select A Device window appears.

12 Click **Save**. The Select Action window appears.

13 Click **Save&Quit**.

14 Restart the network service to start the interface.

```
[admin@platform ~]$ sudo service network restart
```

15 Close the VMware console window.

16 Using an SSH client, log into the Consul VM.

Deploying the Consul VM

Important: You must have access to the Admin node to complete this procedure.

1 On the Consul VM, type the following command to create a **/var/tmp/staging** directory.

```
[admin@consul ~]$ mkdir /var/tmp/staging
```

2 Change to the **/var/tmp/staging** directory.

```
[admin@consul ~]$ cd /var/tmp/staging
```

3 Copy the **cisco-vcs-deployment** zip file from the Admin node to this directory.

Note: Substitute the IP address of your Admin node for the [Admin_IP] entry.

Command Syntax:

```
scp -Crp admin@[Admin_IP]:/opt/cisco/software/admin_node/
cisco-vcs-deployment-*.zip .
```

Example:

```
scp -Crp admin@10.90.44.70:/opt/cisco/software/admin_node/
cisco-vcs-deployment-*.zip .
```

4 Unzip the **cisco-vcs-deployment zip** file and then change to the **scripts** directory.

```
[admin@consul staging]$ unzip cisco-vcs-deployment-*.zip
[admin@consul staging]$ cd cisco-vcs-deployment-*/scripts
```

5 Has a consul encrypt key been generated for your NextX system?

- If **no**, go to the next step.
- If **yes**, go to 8.

6 SSH to the Admin Node and enter the following command to generate a consul encrypt key and record it for future use with the other two Consul VMs you will build.

Important: Only execute this step once as the same key will be used for all nodes in your NextX system.

```
[admin@consul staging]$ ssh admin@[Admin Node IP]
[admin@admin_node ~]$ consul keygen
```

Example Output:

```
qpw3VEZZr4xc5E0bOMObyQ==
```

7 Type **exit** to log out of the Admin Node. You are returned to the Consul node session as admin user in the /var/tmp/staging/cisco-vcs-deployment-[version]/scripts directory.

8 Modify the Consul environment file, **consul.envfile**, in a text editor. A description of each field in the file is shown below.

Notes:

Field	Value
admin_node	Admin Node IP address
consul_datacenter	Set to dc1
consul_servers	List of comma separated Consul server IP addresses
consul_encrypt	Consul encrypt key **Note**: Enter the key generated from Step 6.
hostname	Hostname of the Consul server

```
[admin@consul scripts]$ vi /consul.envfile
```

Example Output:

```
admin_node=172.20.35.10
consul_datacenter=dc1
consul_servers=172.20.33.11,172.20.35.12,172.20.35.13
consul_encrypt=qpw3VEZZr4xc5E0bOMObyQ==
hostname=consul1-prod
```

9 Save and close the file.

10 Enter the following command to deploy the Consul server.

```
[admin@consul scripts]$ sudo ./deploy-consul-server.sh
--envfile=consul.envfile 2>&1 | sudo tee
/var/log/deploy-consul.log
```

Result: The Consul server reboots.

11 Log back into the Consul server as **admin** user.

Transfer X.509 Certificates for TLS Encryption to the Consul VM

Important: The certificate and key pair should have been generated when deploying and configuring the Admin Node. If they have not yet been created, go to the following chapters in the ***Admin Node Installation Guide*** to create them now.

- Chapter 5: Create Environment Files for NextX Nodes
- Chapter 6: Create NextX X.509 Root CA Certificates

This section includes the procedure to transfer the certificate/key pair generated and saved on the Admin Node to the respective node. Once transferred, a procedure is included to verify the configuration of the certificate.

Transferring X.509 Certificates From the Admin Node to the Consul Node

Important: This procedure must be executed for certificates that were generated from an internal root CA or from an external CA.

Complete the following steps to transfer the certificate files created for the Consul nodes from the Admin Node to each respective Consul node.

Note: This procedure is executed on the Admin Node.

1 From a terminal window, log into the Admin Node as **admin** user and then enter the following command to change to **root** user.

```
[admin@adminnode ~]$ sudo -i
```

2 Enter the following command to change to the **/opt/cisco/ca** directory.

```
[root@adminnode ~]# cd /opt/cisco/ca
```

3 Enter the following command and press **Enter** to transfer the appropriate certificate and key pair to the Consul node.

Command Syntax:

```
[root@adminnode ca]# ./manageCerts -P [absolute_path_to_cert]
[absolute_path_to_key] [Consul_IP]
```

Example:

```
[root@adminnode ca]# ./manageCerts -P
/etc/pki/CA/certs/consul.domain.pem
/etc/pki/CA/private/consul.domain.key 10.90.47.246
```

Notes:

- Replace [cert] with the location of the node certificate file (e.g. /etc/pki/CA/certs/[CA.pem]) on the Admin Node.
- Replace [key] with the location of the node private key file (e.g. /etc/pki/CA/private/[CA.key]) on the Admin Node.

- Replace [IP] with the IP address of the Consul node, which is the IP address defined as IP.1 in the [hostname].env file on the Admin Node.

```
./manageCerts -P /etc/pki/CA/certs/consul.default.pem /etc/pki/CA/private/consul.default.key
10.90.47.246
openssl verify -CAfile /etc/pki/CA/cacert.pem /etc/pki/CA/certs/consul.default.pem /etc/pki/CA/
consul.default.pem: OK
openssl verify -CAfile /etc/pki/CA/cacert.pem -purpose sslserver /etc/pki/CA/certs/consul.default
.pem
/etc/pki/CA/certs/consul.default.pem: OK
openssl verify -CAfile /etc/pki/CA/cacert.pem -purpose sslclient /etc/pki/CA/certs/consul.default
.pem
/etc/pki/CA/certs/consul.default.pem: OK
Found X509v3 Extended Key Usage: TLS Web Server Authentication, TLS Web Client Authentication
Found Netscape Cert Type:
Testing connection to 10.90.47.246
ssh -q -t -i /home/admin/.ssh/admin_node_rsa admin@10.90.47.184 sudo mkdir -p "/opt/cisco/ca"
scp -q -i /home/admin/.ssh/admin_node_rsa "/opt/cisco/ca/cascripts.zip" admin@10.90.47.246:
sudo mv -v "cascripts.zip" "/opt/cisco/ca/cascripts.zip"
`cascripts.zip' -> `/opt/cisco/ca/cascripts.zip'
```

4 Were you prompted to verify the SSH RSA key fingerprint:

- If **yes**, type **yes** and press **Enter**. Then go to the next step.
- If **no**, refer to the results section of the next step.

5 When prompted, enter and then re-enter the admin password for the Consul node.

Results:

- The certificate associated private key and truststore are distributed to the Consul node.
- The consul service is restarted.
- A **./manageCerts finished** message displays.

```
Updating permissions for /etc/pki/tls/private/server.key
Updating permissions for /etc/pki/tls/certs/server.crt
Updating permissions for /etc/pki/tls/cacert.pem
Enabling TLS RPC encryption for consul...
Using config file /etc/consul/config.json
Stopping Consul daemon: [ OK ]
Starting Consul daemon: [ OK ]
Nodetype: consul
Please check log file [/var/log/configure_certs20170524.log] for results
/opt/cisco/ca/configure_certs finished
Please check log file [/var/log/manageCerts20170524.log] for results
./manageCerts finished
```

6 Review the logs in the **/var/log** directory.

7 Go to the next section.

Verifying the Consul Certificate Configuration

Complete the following procedure to verify the certificate configuration.

Note: You should still be logged into the Admin Node as the root user in the **/opt/cisco/ca** directory.

1 As **root** user on the Admin Node, enter the following command to check the certificate configuration for the Consul node in which certificates have been generated. A validation of the certificate files occurs.

Command Syntax:

```
./checkConfig -s [hostname].env
```

Example:

```
[root@consul ca]# ./checkConfig -s consul.env
```

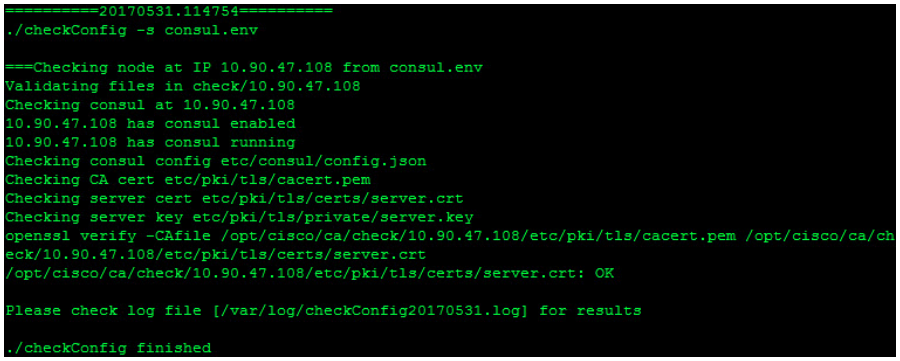

2 Did any errors display?

- If **yes**, review the **/var/log/checkConfig[date].log** file to remedy the issue. Then repeat Step 1. When the issues are corrected, you have completed this procedure.
- If **no**, repeat the previous two steps for each Consul node.

Configuring snmpd on Consul Node

Complete the following procedure to configure snmpd on the Consul node to monitor the Consul service.

1 As **root** user, enter the following command to install the **ServiceManager** package on your Consul node.

```
[root@consul ~]# yum install ServicesManager
```

2 When prompted to confirm the installation, type **y** and press **Enter**. When the installation is finished, a **Complete!** message displays

3 Enter the following command to edit the **/etc/snmp/snmpd.conf** file in a text editor.

```
[root@consul ~]# vi /etc/snmp/snmpd.conf
```

4 Add the following lines to the end of the file:

```
# Monitor consul process and send traps
view    systemview    included   .1.3.6.1.4.1.1429
view    systemview    included   .1.3.6.1.4.1.2021

rocommunity public
rwuser admin
iquerySecName admin
agentSecName admin
proc consul 1 1

monitor -r 5 -u admin -i -o .1.3.6.1.4.1.2021.2.1.2.1 -o
.1.3.6.1.4.1.2021.2.1.101.1 "Consul process has stopped." -o
.1.3.6.1.4.1.2021.2.1.100.1 .1.3.6.1.4.1.2021.2.1.100.1 != 0
monitor -r 5 -u admin -i -o .1.3.6.1.4.1.2021.2.1.2.1 -o
.1.3.6.1.4.1.2021.2.1.101.1 "Consul process is running." -o
.1.3.6.1.4.1.2021.2.1.100.1 .1.3.6.1.4.1.2021.2.1.100.1 == 0

trapsess -v 2c -c public <<VCS_Console_HA_IP>>:162
```

5 Save and close the file.

6 Enter the following command to restart the **snmpd** service.

```
[root@consul ~]# service snmpd restart
```

7 Enter the following command to verify that the **snmpd** is running.

```
[root@consul ~]# service snmpd status
```

Output:

```
snmpd (pid  4609) is running...
```

8 Enter the following command to enable service monitoring.

```
[root@consul ~]#
/opt/cisco/servicesmgr/bin/enableMonitoring.sh
```

9 Enter the following command to verify that the monitoring processes are running.

```
[root@consul ~]# ps -ef | grep -i mon
```

```
dbus      1369     1  0 Jun08 ?        00:00:00 dbus-daemon --system
root      2247     1  0 Jun08 ?        00:15:17 monit
root      5891     1  0 Jun10 ?        00:00:00 /usr/sbin/console-kit-daemon --no-daemon
root     12628 12584  0 16:07 pts/0    00:00:00 grep -i mon
```

10 Have deployed and configured all three Consul nodes?

- If **no**, go back to ***Creating the Consul VM*** (on page 44) and complete all of the procedures up to and including this section.
- If **yes**, go to ***Verifying Consul Functionality*** (on page 53).

Verifying Consul Functionality

Important: All three Consul servers must be deployed and installed to successfully execute the scripts within this procedure.

Complete the following steps to verify the functionality of the three Consul servers.

1 As **admin** user, enter the following command to verify that the consul nodes joined a cluster and that the Consul service is synced.

```
[admin@consul ~]$ consul monitor
```

Example Output:

```
2017/06/28 14:00:03 [INFO] raft: Restored from snapshot 530-1720435-1498648183536
2017/06/28 14:00:03 [INFO] raft: Initial configuration (index=653): [{Suffrage:Voter ID:10.
90.47.108:8300 Address:10.90.47.108:8300} {Suffrage:Voter ID:10.90.47.109:8300 Address:10.9
0.47.109:8300} {Suffrage:Voter ID:10.90.47.177:8300 Address:10.90.47.177:8300}]
2017/06/28 14:00:03 [INFO] raft: Node at 10.90.47.177:8300 [Follower] entering Follower sta
te (Leader: "")
2017/06/28 14:00:03 [INFO] serf: EventMemberJoin: consul3 10.90.47.177
2017/06/28 14:00:03 [INFO] consul: Adding LAN server consul3 (Addr: tcp/10.90.47.177:8300)
(DC: dc1)
2017/06/28 14:00:03 [INFO] consul: Raft data found, disabling bootstrap mode
2017/06/28 14:00:03 [INFO] serf: EventMemberJoin: vodwater 10.90.45.181
2017/06/28 14:00:03 [INFO] serf: EventMemberJoin: consul 10.90.47.108
.
.
2017/06/28 14:00:03 [INFO] serf: EventMemberJoin: consul2 10.90.47.109
2017/06/28 14:00:03 [INFO] serf: EventMemberJoin: vodwaterDtacs 10.90.47.104
2017/06/28 14:00:03 [INFO] consul: Adding LAN server consul (Addr: tcp/10.90.47.108:8300) (
DC: dc1)
2017/06/28 14:00:03 [INFO] consul: Adding LAN server consul2 (Addr: tcp/10.90.47.109:8300)
(DC: dc1)
2017/06/28 14:00:03 [INFO] serf: EventMemberJoin: consul3.dc1 10.90.47.177
2017/06/28 14:00:03 [WARN] serf: Failed to re-join any previously known node
2017/06/28 14:00:03 [INFO] consul: Adding WAN server consul3.dc1 (Addr: tcp/10.90.47.177:83
00) (DC: dc1)
2017/06/28 14:00:03 [INFO] agent: Joining cluster...
2017/06/28 14:00:03 [INFO] agent: (LAN) joining: [10.90.47.177 10.90.47.109 10.90.47.108]
2017/06/28 14:00:03 [INFO] agent: (LAN) joined: 3 Err: <nil>
2017/06/28 14:00:03 [INFO] agent: Join completed. Synced with 3 initial agents
2017/06/28 14:00:04 [INFO] agent: Synced service 'consul'
2017/06/28 14:00:23 [INFO] agent.rpc: Accepted client: 127.0.0.1:41106
```

2 Enter the following commands to verity the the Consul server is running and that you are a member of the cluster.

```
[admin@consul ~]$ sudo lsof -Pni :8500
```

```
COMMAND  PID   USER   FD   TYPE  DEVICE SIZE/OFF NODE NAME
consul  5194 consul   15u  IPv4 8307854      0t0  TCP 127.0.0.1:8500 (LISTEN)
```

```
[admin@consul ~]$ consul members
```

```
Node                  Address             Status  Type    Build  Protocol  DC
ec80                  10.90.178.166:8301  alive   client  0.7.2  2         dc1
consul                10.90.47.108:8301   alive   server  0.7.2  2         dc1
consul2               10.90.47.109:8301   alive   server  0.7.2  2         dc1
consul3               10.90.47.177:8301   alive   server  0.7.2  2         dc1
vodwater              10.90.45.181:8301   alive   client  0.7.2  2         dc1
vodwaterDtacs         10.90.47.104:8301   alive   client  0.7.2  2         dc1
```

Install the VCS Console

The Video Controller Suite (VCS) Console is the Web UI infrastructure that allows you to manage multi-server nodes in your video system (e.g. ECs, DTACS). Your NextX environment will include two VCS Console nodes to create a high availability (HA) environment.

Please ensure that your system meets the following requirements before proceeding.

Important: If any of these requirements have not been completed, please do so now.

- Oracle RAC installed
- ECS database users successfully created
- Linux platform template built
- Admin Node deployed and running
- NextX X.509 Root Certificate Authority (CA) Certificates are created on the Admin Node
- VCS Console bundle is deployed on the Admin node
- Consul node deployed and running
- Three IP addresses are required for the HA environment
  - VCS Console 1 IP
  - VCS Console 2 IP
  - VCS HA Virtual IP (VIP)

Creating the VCS Console VM

Important: If you are using vSphere client to deploy virtual machines, you cannot create the VCS Console VM using a template. Refer to ***Procedures When Using vSphere Client*** (on page 221).

Complete the following procedures to deploy the VCS Console VM from the vSphere Web UI.

1. From the vCenter Web UI, click **VMs and Templates**.
2. Locate and select the CSCOlxplat template that was built using the ***Admin Node Installation Guide***.
3. Right-click the template and select **Deploy VM from this Template**. The Deploy From Template window opens.
4. In the text box, enter a name for the VM you are creating and then select the datacenter where it will be deployed. Click **Next.**

5 Select the appropriate ESXi host where you want to deploy the VM and click **Next**.

6 From the **Select virtual disk format** dropdown menu, maintain the **Same format as source** default. Then ensure that the appropriate datastore is selected.

7 Click **Next** and then click **Next** again.

8 Review the settings and click **Finish**.

Reconfiguring the VCS Console VM

Important: If you are using vSphere client to deploy and configure VMs, refer to *Reconfiguring the Virtual Network Using vSphere Client* (on page 223).

Complete the following procedure to reconfigure the virtual hardware on the VCS Console node.

1 Locate and select the VCS Console VM.

2 Right-click the VM and select **Edit Settings**. The Edit Settings window appears.

3 From the **CPU** text box, change the value to **8**.

4 From the **Memory** text box, change the value to **16** and then click the dropdown box next to the value and select **GB**.

5 From the **Hard disk 1** text box, change the value to **160 GB**.

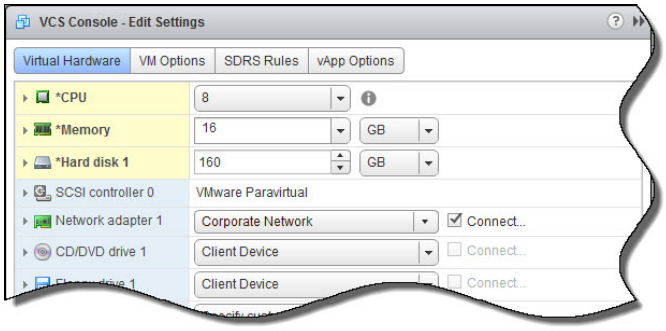

6 Is the Oracle RAC set up on a different network?

- If **yes**, go to step 7.
- If **no**, go to Step 11.

7 From the **New device** dropdown menu, select **Network** and then click **Add**.

8 Click the arrow next to the **New Network** entry to view the details.

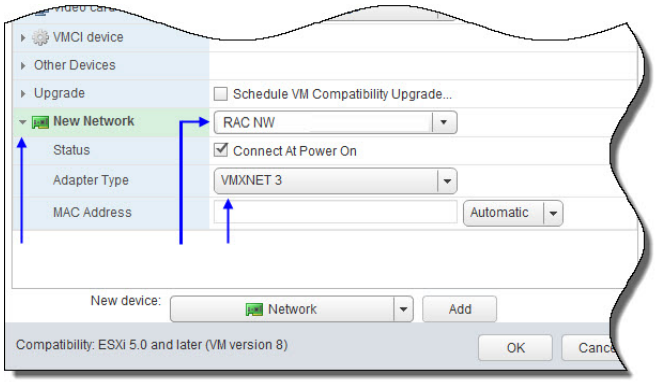

9 From the **New Network** dropdown menu, select the network associated with the RAC database.

10 From the **Adapter Type** dropdown menu, select **VMXNET 3**.

11 Click **OK**. The VM is reconfigured.

Configuring the VCS Console Network Interface With a Static IP

1 Select and right-click the VCS Console VM and select **Power On**.

2 Select and right-click the VCS Console VM again and select **Open Console**. A VMware console window opens in a new tab.

3 Log into the VM with the following credentials.

Username: admin

Password: [password created for Linux platform template]

4 Type the following command line utility to configure the network and the DNS settings. The Select Action window appears.

```
[admin@platform ~]$ sudo system-config-network
```

5 Select **Device configuration** and press **Enter**.

6 Highlight **eth0** and press **Enter**.

7 Press the **Tab** key until you **Use DHCP i**s highlighted. Then press the **Spacebar** to unselect this option.

8 Tab to each field to enter the following values.

Note: DNS entries are optional.

- **Static IP**
- **Netmask**
- **Default gateway IP**
- **Primary DNS Server**
- **Secondary DNS Server**

9 Verify that **Peer DNS** is selected.

10 Press the **Tab** key until you **Controlled by NetworkManager i**s highlighted. Then press the **Spacebar** to unselect this option.

11 Press **Tab** to highlight **Ok** and press **Enter**. The Select A Device window appears.

12 Click **Save**. The Select Action window appears.

13 Click **Save&Quit**.

14 Restart the network service to start the interface.

```
[admin@platform ~]$ sudo service network restart
```

15 Is the RAC in a different subnet than the VM?

- If **no**, you have completed this procedure.
- If **yes**, repeat the above procedure to configure a secondary interface on **eth1**.

16 Verify that you can ping the RAC database.

17 Did you successfully ping the RAC database?

- If **yes**, you have completed this procedure.
- If **no**, troubleshoot your network or call Cisco Services.

Deploying the VCS Console VM

1 Using an SSH client, log into the VCS Console VM.

2 On the VCS Console VM, type the following command to create a **/var/tmp/staging** directory.

```
[admin@vcsconsole ~]$ mkdir /var/tmp/staging
```

3 Change to the **/var/tmp/staging** directory.

```
[admin@vcsconsole ~]$ cd /var/tmp/staging
```

4 Copy the **cisco-vcs-deployment** zip file from the Admin node to this directory.

Command Syntax:

```
scp -Crp admin@[Admin_IP]:/opt/cisco/software/admin_node/
cisco-vcs-deployment-*.zip .
```

Example:

```
scp -Crp admin@10.90.44.70:/opt/cisco/software/admin_node/
cisco-vcs-deployment-*.zip .
```

5 Unzip the cisco-vcs-deployment file and then change to the **scripts** directory.

```
[admin@vcsconsole staging]$ unzip cisco-vcs-deployment-*.zip
[admin@vcsconsole staging]$ cd cisco-vcs-deployment-*/scripts
```

6 Modify the VCS Console environment file, **vcsconsole.envfile**. A description of each field in the file is shown below.

Field	Value
install_labcerts	Set to false
persistence_port	RAC database listening port
persistence_type	Database engine name
persistence_name	Database name
consul_datacenter	Consul datacenter name
admin_node	Admin node name or IP address
persistence_host	RAC scan IP address
persistence_user	RAC VCS Console database user name
persistence_password	RAC VCS Console database user password
consul servers	All Consul server IP addresses separated by a comma (no spaces)
consul_encrypt	Consul encryption key
hostname	VCS Console hostname

Default vcsconsole.envfile:

```
install_labcerts=false
persistence_port=1535
persistence_type=oracle
persistence_name=CABHE
consul_datacenter=dc1
admin_node=
persistence_host=
persistence_user=
persistence_password=
consul_servers=
consul_encrypt=
```

Example:

```
install_labcerts=false
persistence_port=1535
persistence_type=oracle
persistence_name=CABHE
consul_datacenter=dc1
admin_node=10.90.44.70
persistence_host=172.20.36.5
persistence_user=VCSPROD_VCSCONSOLE
persistence_password=vcsprod
consul_servers=172.20.35.11,172.20.35.12,172.20.35.13
consul_encrypt=qpw3VEZZr4xc5E0b0MObyQ==
hostname=vcs-prod
```

7 Save and close the file.

8 Enter the following command to execute the VCS Console deploy script.

```
[admin@vcsconsole scripts]$ sudo ./deploy-vcsconsole.sh
--envfile=vcsconsole.envfile 2>&1 | sudo tee
/var/log/deploy-vcsconsole.log
```

Results: When the script completes, the VCS Console VM reboots.

Transfer X.509 Certificates for TLS Encryption to the VCS Console Nodes

Important: The certificate and key pair should have been generated when deploying and configuring the Admin Node. If they have not yet been created, go to the following chapters in the ***Admin Node Installation Guide*** to create them now.

- Chapter 5: Create Environment Files for NextX Nodes
- Chapter 6: Create NextX X.509 Root CA Certificates

This section includes the procedure to transfer the certificate/key pair generated and saved on the Admin Node to the respective node. Once transferred, a procedure is included to verify the configuration of the certificate.

Transferring X.509 Certificates From the Admin Node to the VCS Console Node

Important: This procedure must be executed for certificates that were generated from an internal root CA or from an external CA.

Complete the following steps to transfer the certificate files created for the VCS Console nodes from the Admin Node to each respective VCS Console node.

Note: This procedure is executed on the Admin Node.

1 From a terminal window, log into the Admin Node as **admin** user and then enter the following command to change to **root** user.

```
[admin@adminnode ~]$ sudo -i
```

2 Enter the following command to change to the **/opt/cisco/ca** directory.

```
[root@adminnode ~]# cd /opt/cisco/ca
```

3 Enter the following command and press **Enter** to transfer the appropriate certificate and key pair to the VCS Console node.

Command Syntax:

```
[root@adminnode ca]# ./manageCerts -P [absolute_path_to_cert]
[absolute_path_to_key] [VCS Console_IP]
```

Example:

```
[root@adminnode ca]# ./manageCerts -P
/etc/pki/CA/certs/vcsconsole.domain.pem
/etc/pki/CA/private/vcsconsole.domain.key 10.90.47.246
```

Notes:

- Replace [cert] with the location of the node certificate file (e.g. /etc/pki/CA/certs/[CA.pem]) on the Admin Node.
- Replace [key] with the location of the node private key file (e.g. /etc/pki/CA/private/[CA.key]) on the Admin Node.

- Replace [IP] with the IP address of the VCS Console node, which is the IP address defined as IP.1 in the [hostname].env file on the Admin Node.

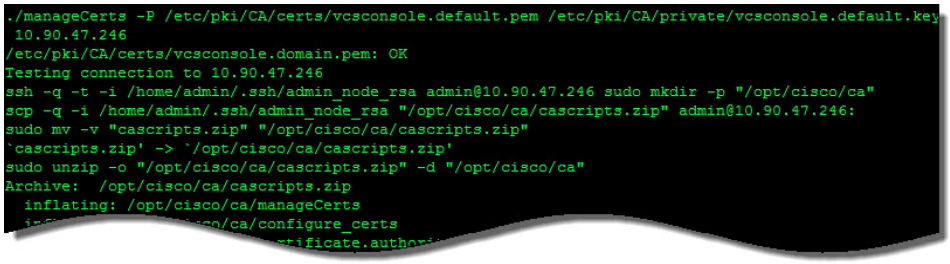

4 Where you prompted to verify the SSH RSA key fingerprint?

- If **yes**, type **yes** and press **Enter**. Then go to the next step.
- If **no**, refer to the results section of the next step.

5 When prompted, enter and then re-enter the admin password for the VCS Console node.

Results:

- The certificate associated private key and truststore are distributed to the VCS Console node.
- The VCS Console security configuration is updated.
- The VCS Console service is restarted.
- A **./manageCerts finished** message displays.

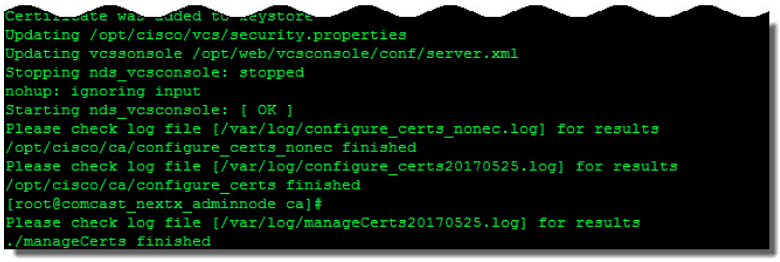

6 Review the logs in the **/var/log** directory.

7 Go to the next section.

Verifying the VCS Console Certificate Configuration

Once the VCS Console node is configured, you can execute the **checkConfig** script to verify the configuration. Complete the following steps to verify these parameters.

Note: You should still be logged into the Admin Node as **admin** user in the /opt/cisco/ca directory.

1 As **root** user on the Admin Node, enter the following command to check the certificate configuration for each of the VCS Console nodes in which certificates have been generated. A validation of the certificate files occurs.

Command Syntax:

```
sudo ./checkConfig -s [hostname].env
```

Example:

```
[admin@adminnode ca]$ sudo ./checkConfig -s vcsconsole1.env
```

```
==========20170531.120209==========
./checkConfig -s vcsconsole1.env

===Checking node at IP 10.90.47.145 from vcsconsole1.env
Validating files in check/10.90.47.145
Checking vcs console at 10.90.47.145
10.90.47.145 has consul enabled
10.90.47.145 has consul running
10.90.47.145 has vcsconsole enabled
10.90.47.145 has vcsconsole running
10.90.47.145 has Alarm_Manager enabled
10.90.47.145 Alarm_Manager is running
Checking consul config etc/consul/config.json
Checking CA cert etc/pki/tls/cacert.pem
Checking server cert etc/pki/tls/certs/server.crt
Checking server key etc/pki/tls/private/server.key
Checking security properties opt/cisco/vcs/security.properties
Checking VCS console config opt/web/vcsconsole/conf/server.xml
openssl verify -CAfile /opt/cisco/ca/check/10.90.47.145/etc/pki/tls/cacert.pem /opt/cisco/ca/check
/10.90.47.145/etc/pki/tls/certs/server.crt
/opt/cisco/ca/check/10.90.47.145/etc/pki/tls/certs/server.crt: OK

Please check log file [/var/log/checkConfig20170531.log] for results

./checkConfig finished
```

2 Did any errors display?

- If **yes**, review the **/var/log/checkConfig[date].log** file to remedy the issue. Then repeat Step 1. When the issues are corrected, you have completed this procedure.
- If **no**, go to *Configuring VCS Console to Use Consul for Directory Services* (on page 61).

Configuring VCS Console to Use Consul for Directory Services

1 As **root** user on the VCS Console, enter the following command to configure the VCS Console to use the Consul node for its directory services.

Field	Definition
-sdClientType consul	Service directory client
-sdhost localhost	Localhost where the Consul local agent is running
sdport 8500	Consul listening port
-sdConsulLocalAgent true	Set to "true" for Consul
-sdConsulLocalAgentPort	Set to 8500 if the Consul agent is running locally

Command Syntax:

```
/opt/vcs/bin/vcsutils.sh -updateSDProperties -sdClientType
consul -sdhost localhost -sdDataCenter dc1 -sdport 8500
-sdConsulLocalAgent true -sdConsulLocalAgentPort 8500
```

Result: The file /opt/vcs/conf/sdconfig.properties is updated.

2 Enter the following command to review the **sdconfig.properties** file.

```
[root@vcsconsole~]# cat /opt/vcs/conf/sdconfig.properties
```

Install and Configure Alarm Manager

Complete the following steps to install and then configure Alarm Manager on the VCS Console node.

1 Enter the following command to stop the Tomcat service.

Note: This results of this command may show a Connection Refused error message if Tomcat is not running. This is a normal message and can be ignored.

```
[root@vcsconsole ~]# /opt/apache/tomcat70/bin/shutdown.sh
```

2 Enter the following command to install Alarm Manager.

```
[root@vcsconsole~]# yum install Alarm_Manager
```

3 When prompted to confirm the installation, type **y**.

4 When the install completes, edit the following file to configure Alarm Manager database access.

```
[root@vcsconsole~]# vi /opt/jdbc.properties
```

5 Update only the following fields to reflect your system environment.

Field	Definition
password	Password for the ALARMS database user
url	RAC scan IP address or hostname, the port to the Oracle listening port (1535) and the Oracle Site Identifiier (SID) which will be CABHE
username	Username for the ALARMS database user
dbport	1535
dbhost	IP address of the RAC scan IP address

Example:

```
connectionPoolSize=5
testOnBorrow=true
validationQuery=select 1 from dual
password=vcsprod
driverClassName=oracle.jdbc.driver.OracleDriver
url=jdbc:oracle:thin:@172.20.36.5:1535/CABHE
username=VCSPROD_ALARMS
dbport=1535
dbhost=172.20.36.5
autoPwRefresh=false
xmp.persistence.tqexclusion.file=/WEB-INF/config/TqNotRequired.txt
```

6 Save and close the file.

7 Modify **/opt/apache/tomcat70/conf/server.xml** to set the following values to **true**.

```
unpackWARs="true" autoDeploy="true"
```

8 Save and close the file.

9 Enter the following command to start Alarm Manager.

```
[root@vcsconsole ~]# service Alarm_Manager start
```

10 Enter the following command to verify that Alarm Manager and the Decap processes started successfully.

```
[root@vcsconsole ~]# service Alarm_Manager status
```

Result: An Alarm Manager and a DECAP process is running message appears along with its PID numbers.

Example:

```
[admin@vcsconsole ~]$ sudo service Alarm_Manager status
Alarm_Manager process(es) running (15292)
Decap process(es) running (15364 15369 )
```

11 Enter the following command to verify that Alarm Manager started successfully.

```
[root@vcsconsole ~]# less
/opt/apache/tomcat70/logs/catalina.out
```

12 Enter the following command to install the Alarm Manager Web UI.

```
[root@vcsconsole ~]# yum install AlarmManagement_UI
```

13 When prompted, enter **y** to confirm the installation.

14 After the package installs, enter the following command to edit the Alarms Manager Web UI configuration file.

```
[root@vcsconsole ~]# vi /opt/alarm_dao.properties
```

15 Update the following fields to reflect your system environment.

Field	Definition
db.url=	RAC scan IP address or hostname, the port to the Oracle listening port (1535) and the Oracle Site Identifiier (SID) which will be CABHE
db.username=	Username for the ALARMS database user
db.password=	Username for the ALARMS database user
hibernate.default_schema=	Username for the ALARMS database user

Example:

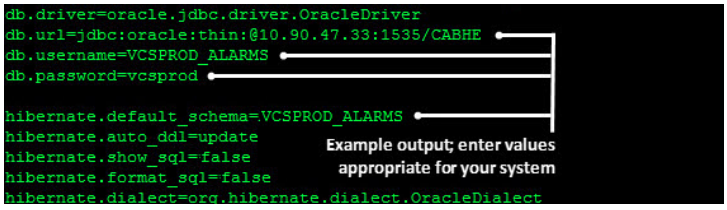

Starting the VCS Console

Complete the following procedure as **root** user on the VCS Console.

1 Enter the following command to verify that the VCS Console is a member of the Consul cluster.

```
[root@vcsconsole ~]# consul members
```

Example:

```
Node            Address              Status  Type    Build  Protocol  DC
consul1-prod  172.20.35.11:8301      alive   server  0.7.0  2         dc1
consul2-prod  172.20.35.12:8301      alive   server  0.7.0  2         dc1
consul3-prod  172.20.35.13:8301      alive   server  0.7.0  2         dc1
vcs1-prod     172.20.35.18:8301      alive   client  0.7.0  2         dc1
```

2 Enter the following command to enable the VCS Console to start on reboot.

```
[root@vcsconsole ~]# chkconfig vcsconsole on
```

3 Enter the following command to restart the VCS Console.

```
[root@vcsconsole ~]# service vcsconsole restart
```

```
[root@situcs5-vcs-vcsconsole bin]# service vcsconsole start
Starting nds_vcsconsole:                                   [  OK  ]
```

4 From a supported Web browser, enter the following command to access the VCS Console Web UI.

URL Format:

```
https://[VCS Console VIP]:6605/vcsconsole
```

Example:

```
https//10.90.47.33:6605/vcsconsole
```

5 When prompted to login in, enter the following default login credentials. The VCS Console Home page displays. You are prompted to change the password.

User Name: root

Password: Public123

6 When prompted, enter a new password.

7 When prompted to re-enter the password, enter the new password again. The VCS Console Home page displays.

Example Output:

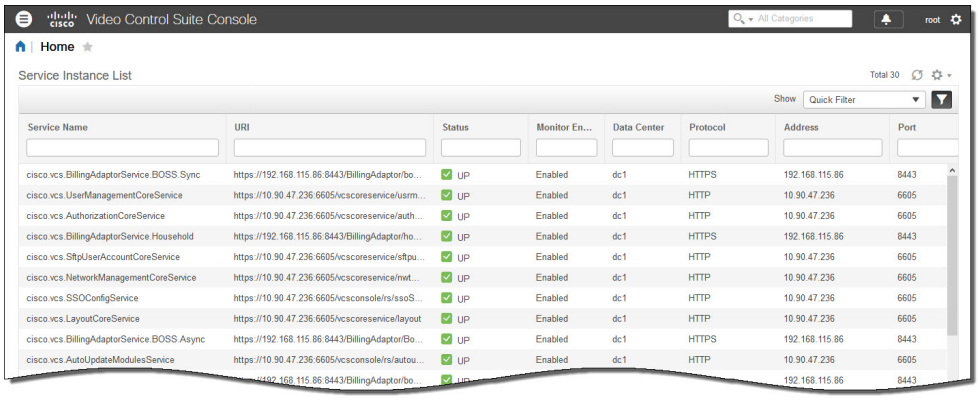

8 Were you able to access the VCS Console Web UI?

- If **yes**, you have completed this procedure.
- If **no**, go to the next step.

9 From your Firefox Web browser, click the **Open** menu icon, ☰, and select **Options**.

10 Click **Advanced** and and then click the **Certifications** tab.

11 Deselect the **Query OCSP responder servers to confirm the current validity of the certificates** box.

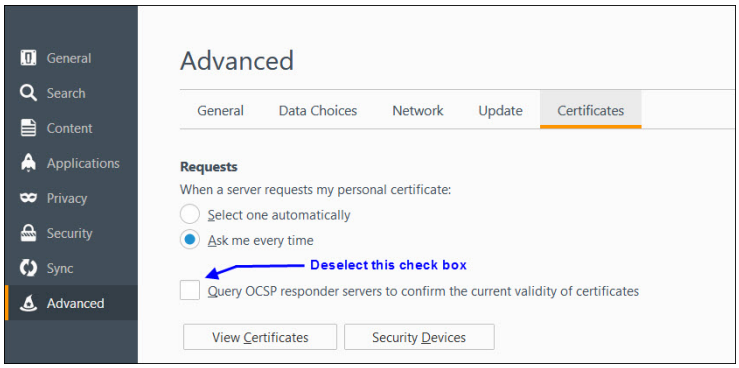

12 Repeat Steps 4 through 6 to access the VCS Console Web UI.

Configuring snmpd on the VCS Console Node

Complete the following procedure to configure SNMPd on the VCS Console node to monitor the following services:

- Consul
- AlarmManager - Tomcat
- VCS Console - Tomcat

1 Enter the following command to install the **ServiceManager** package on your VCS Console node.

```
[root@vcsconsole ~]# yum install ServicesManager
```

2 When prompted to confirm the installation, type **y** and press **Enter**. When the installation is finished, a **Complete!** message displays.

3 As **root** user, enter the following command to edit the **/etc/snmp/snmpd.conf** file in a text editor.

```
[root@vcsconsole ~]# vi /etc/snmp/snmpd.conf
```

4 Add the following lines to the end of the file:

```
# Monitor consul process and send traps
view    systemview    included   .1.3.6.1.4.1.1429
view    systemview    included   .1.3.6.1.4.1.2021

rocommunity public
rwuser admin
iquerySecName admin
agentSecName admin
proc consul 1 1
proc alarmsvcsmon 1 1
proc vcsconsvcsmon 1 1

monitor -r 5 -u admin -i -o .1.3.6.1.4.1.2021.2.1.2.1 -o
.1.3.6.1.4.1.2021.2.1.101.1 "Consul process has stopped." -o
.1.3.6.1.4.1.2021.2.1.100.1 .1.3.6.1.4.1.2021.2.1.100.1 != 0
monitor -r 5 -u admin -i -o .1.3.6.1.4.1.2021.2.1.2.1 -o
.1.3.6.1.4.1.2021.2.1.101.1 "Consul process is running." -o
.1.3.6.1.4.1.2021.2.1.100.1 .1.3.6.1.4.1.2021.2.1.100.1 == 0

monitor -r 5 -u admin -i -o .1.3.6.1.4.1.2021.2.1.2.2 -o
.1.3.6.1.4.1.2021.2.1.101.2 "ECS Alarm Mgr process stopped." -
o .1.3.6.1.4.1.2021.2.1.100.2 .1.3.6.1.4.1.2021.2.1.100.2 != 0
monitor -r 5 -u admin -i -o .1.3.6.1.4.1.2021.2.1.2.2 -o
.1.3.6.1.4.1.2021.2.1.101.2 "ECS Alarm Mgr process running." -
o .1.3.6.1.4.1.2021.2.1.100.2 .1.3.6.1.4.1.2021.2.1.100.2 == 0

monitor -r 5 -u admin -i -o .1.3.6.1.4.1.2021.2.1.2.3 -o
.1.3.6.1.4.1.2021.2.1.101.3 "VCS Console processes stopped." -
o .1.3.6.1.4.1.2021.2.1.100.3 .1.3.6.1.4.1.2021.2.1.100.3 != 0
monitor -r 5 -u admin -i -o .1.3.6.1.4.1.2021.2.1.2.3 -o
.1.3.6.1.4.1.2021.2.1.101.3 "VCS Console processes running." -
o .1.3.6.1.4.1.2021.2.1.100.3 .1.3.6.1.4.1.2021.2.1.100.3 == 0

trapsess -v 2c -c public <<VCS Console HA IP>>:162
```

5 Save and close the file.

6 Enter the following command to restart the **snmpd** service.

```
[root@vcsconsole ~]# service snmpd restart
```

7 Enter the following command to verify that the **snmpd** is running.

```
[root@vcsconsole ~]# service snmpd status
```

Output:

```
snmpd (pid  4609) is running...
```

8 Enter the following command to enable service monitoring.

```
[root@vcsconsole ~]#
/opt/cisco/servicesmgr/bin/enableMonitoring.sh
```

9 Enter the following command to verify that the monitoring processes are running.

```
[root@vcsconsole ~]# ps -ef | grep -i mon
```

```
dbus      1364     1  0 Jul14 ?        00:00:00 dbus-daemon --system
root      2060     1  0 Jul14 ?        00:04:59 monit
root     18700  2036  0 12:16 ?        00:00:00 /opt/cisco/servicesmgr/bin/alarmsvcsmon
root     18706  1987  0 12:16 ?        00:00:00 /opt/cisco/servicesmgr/bin/vcsconsvcsmon
root     20454 20439  0 13:40 pts/0    00:00:00 grep -i mon
```

Configure VCS Console High Availability (HA)

VCS Console HA is handled by the following two open source software applications.

- **keepalived**
- **HAProxy**

The keepalived application manages the Virtual IP (VIP) and maintains the VIP on the appropriate server. Initially, the VIP is configured on the Master server and is then transferred to the Slave server only when the Master VCS Console service is not responding. If the Master server is down and becomes healthy again, the VIP is transferred back to the Master server.

HAProxy provides high availability and load balancing capabilities. HAproxy handles the VCS Console Web UI requests and distributes them using a roundrobin mechanism. HAProxy also has rules (sticky rules) to keep a session, coming from a client, established to the same VCS Console server.

The following diagram illustrates the HA environment.

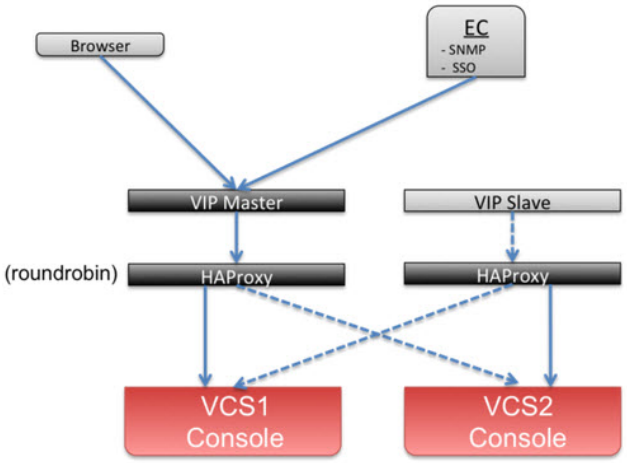

Build a Secondary VCS Console

To create a secondary VCS Console server, go to the ***Install the VCS Console*** (on page 54) and follow all of the procedures until you return to this section. Use a different hostname (i.e. vcsconsole2) and IP address for the secondary VCS Console.

Once both VCS Consoles are installed and configured, go to the next section, ***Installing HAProxy on Each VCS Console***.

Installing HAProxy on Each VCS Console VM

Important: To complete this procedure, both VCS Consoles should be deployed and configured.

Complete the following procedure to install HAProxy on each VCS Console VM. In these procedure, the primary VCS Console is the Master server and the secondary VCS Console is the Slave server.

1. From the primary VCS Console VM, enter the following command to install HAProxy on the **Master** server.

   ```
   [root@vcsconsole ~]# yum install haproxy
   ```

2. When prompted to confirm the installation, enter **y**.
3. When the installation completes, repeat Steps 1 through 2 on the **Slave** server.
4. On the **Master** server, make a backup of the haproxy.cfg file.

   ```
   [root@vcsconsole ~]# mv /etc/haproxy/haproxy.cfg
   /etc/haproxy/haproxy.cfg.def
   ```

5. Open a new file named **haproxy.cfg** in a text editor.

   ```
   [root@vcsconsole ~]# vi /etc/haproxy/haproxy.cfg
   ```

6. Copy and paste the following content into the **haproxy** file.

   ```
   #---------------------------------------------------------
   # Example configuration for a possible web application.  See
   the
   # full configuration options online.
   #
   #   http://haproxy.1wt.eu/download/1.4/doc/configuration.txt
   #
   #---------------------------------------------------------
   #[MASTER]
   #---------------------------------------------------------
   # Global settings
   #---------------------------------------------------------
   global
       # to have these messages end up in /var/log/haproxy.log
   you will
       # need to:
       #
   ```

```
    # 1) configure syslog to accept network log events.  This
is done
    #    by adding the '-r' option to the SYSLOGD_OPTIONS in
    #    /etc/sysconfig/syslog
    #
    # 2) configure local2 events to go to the
/var/log/haproxy.log
    #   file. A line like the following can be added to
    #   /etc/sysconfig/syslog
    #
    #    local2.*                       /var/log/haproxy.log
    #
    log         127.0.0.1 local2
    maxconn     40000
    daemon


#---------------------------------------------------------
# common defaults that all the 'listen' and 'backend' sections
will
# use if not designated in their block
#---------------------------------------------------------
defaults
    mode                    http
    log                     global
    option                  dontlognull
    option                  redispatch
    timeout queue           1000s
    timeout connect         5000
    timeout client          86400000
    timeout server          86400000


#---------------------------------------------------------
# main frontend which proxys to the backends
#---------------------------------------------------------
frontend  https_frontend
    bind *:443
    mode tcp
    default_backend             app


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

```
# static backend for serving up images, stylesheets and such
#--------------------------------------------------------------


#--------------------------------------------------------------
# round robin balancing between the various backends
#--------------------------------------------------------------
backend app
    balance     roundrobin


mode tcp
    stick-table type binary len 32 size 30k expire 30m
    acl clienthello req_ssl_hello_type 1
    acl serverhello rep_ssl_hello_type 2
    tcp-request inspect-delay 5s
    tcp-request content accept if clienthello


    tcp-response content accept if serverhello
    stick on payload_lv(43,1) if clienthello
    stick store-response payload_lv(43,1) if serverhello
    server s1 10.90.47.21:6605 check
    server s2 10.90.47.111:6605 check


# To watch report on browser [FOR MORE INFO
http://tecadmin.net/how-to-configure-haproxy-statics/]
listen  stats   *:1936
        mode            http
        log             global
        maxconn 10
        timeout client      100s
        timeout server      100s
        timeout connect      100s
        timeout queue   100s
        stats enable
        stats hide-version
        stats refresh 30s
        stats show-node
        stats auth admin:password
        stats uri  /haproxy?stats
```

7 Within the content you pasted, customize the following fields for your VCS HA solution.

- For the **server s1** entry, replace the IP with the Master IP address
- For the **server s2** entry, replace the IP with the Slave IP address
- For the **stats auth** entry, type the VCS Console VM login credentials

Example:

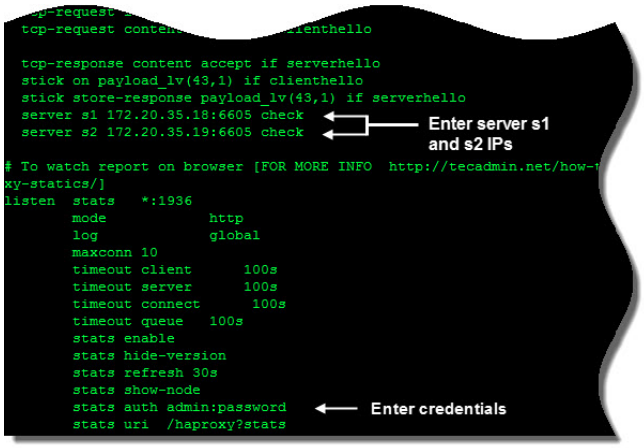

8 Save and close the file.

9 On the **Slave** server, enter the following command to create a backup of the haproxy.cfg file

```
[root@vcsconsole2 ~]# mv /etc/haproxy/haproxy.cfg
/etc/haproxy/haproxy.cfg.def
```

10 Open a new file named **haproxy.cfg** in a text editor.

```
[root@vcsconsole2 ~]# vi /etc/haproxy/haproxy.cfg
```

11 Copy and paste the following content into the **haproxy** file.

```
#---------------------------------------------------------------
# Example configuration for a possible web application.  See
the
# full configuration options online.
#
#   http://haproxy.1wt.eu/download/1.4/doc/configuration.txt
#
#---------------------------------------------------------------
#[SLAVE]
#---------------------------------------------------------------
# Global settings
#---------------------------------------------------------------
global
    # to have these messages end up in /var/log/haproxy.log
you will
```

```
    # need to:
    #
    # 1) configure syslog to accept network log events.  This
is done
    #    by adding the '-r' option to the SYSLOGD_OPTIONS in
    #    /etc/sysconfig/syslog
    #
    # 2) configure local2 events to go to the
/var/log/haproxy.log
    #   file. A line like the following can be added to
    #   /etc/sysconfig/syslog
    #
    #    local2.*                       /var/log/haproxy.log
    #
    log         127.0.0.1 local2
    maxconn     40000
    daemon


#---------------------------------------------------------------------
# common defaults that all the 'listen' and 'backend' sections
will
# use if not designated in their block
#---------------------------------------------------------------------
defaults
    mode                    http
    log                     global
    option                  dontlognull
    option                  redispatch
    timeout queue           1000s
    timeout connect         5000
    timeout client          86400000
    timeout server          86400000


#---------------------------------------------------------------------
# main frontend which proxys to the backends
#---------------------------------------------------------------------
frontend  https_frontend
    bind *:443
    mode tcp
    default_backend             app
```

```
#---------------------------------------------------------------------
# static backend for serving up images, stylesheets and such
#---------------------------------------------------------------------

#---------------------------------------------------------------------
# round robin balancing between the various backends
#---------------------------------------------------------------------
backend app
    balance     roundrobin

mode tcp
    stick-table type binary len 32 size 30k expire 30m
    acl clienthello req_ssl_hello_type 1
    acl serverhello rep_ssl_hello_type 2
    tcp-request inspect-delay 5s
    tcp-request content accept if clienthello

    tcp-response content accept if serverhello
    stick on payload_lv(43,1) if clienthello
    stick store-response payload_lv(43,1) if serverhello
    server s1 10.90.47.21:6605 check
    server s2 10.90.47.111:6605 check

# To watch report on browser [FOR MORE INFO
http://tecadmin.net/how-to-configure-haproxy-statics/]
listen  stats   *:1936
        mode            http
        log             global
        maxconn 10
        timeout client      100s
        timeout server      100s
        timeout connect      100s
        timeout queue   100s
        stats enable
        stats hide-version
        stats refresh 30s
        stats show-node
        stats auth admin:password
        stats uri  /haproxy?stats
```

12 Within the content you pasted, customize the following fields or your VCS HA solution.

- For the **server s1** entry, replace the IP with the Master IP address
- For the **server s2** entry, replace the IP with the Slave IP address
- For the **stats auth** entry, type the VCS Console VM login credentials.

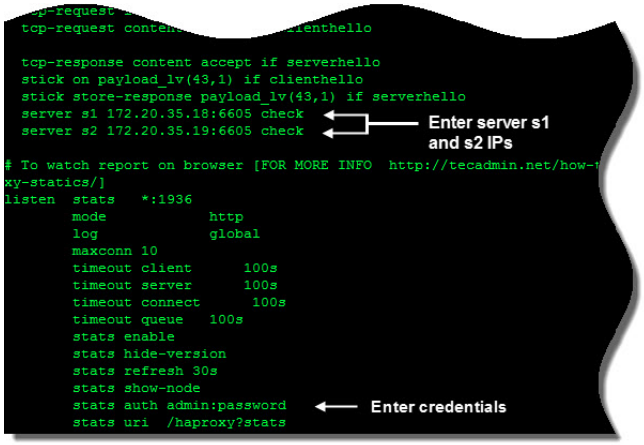

13 Save and close the file.

14 On the **Master** server, enter the following command to configure HAProxy to start when booted up.

```
[root@vcsconsole ~]# chkconfig haproxy on
```

15 On the **Master** server, enter the following command to start HAProxy.

```
[root@vcsconsole ~]# service haproxy start
```

16 Repeat Steps 14 through 15 on the **Slave** server.

17 From a Web browser, enter the following command to verify the status of HAProxy on the **Master** server.

Note: To log into the server, enter the credentials defined in the respective haproxy.cfg file.

```
http://[Master IP Address]:1936/haproxy?stats
```

Result: Both the s1 and s2 VCS Consoles rows should be green.

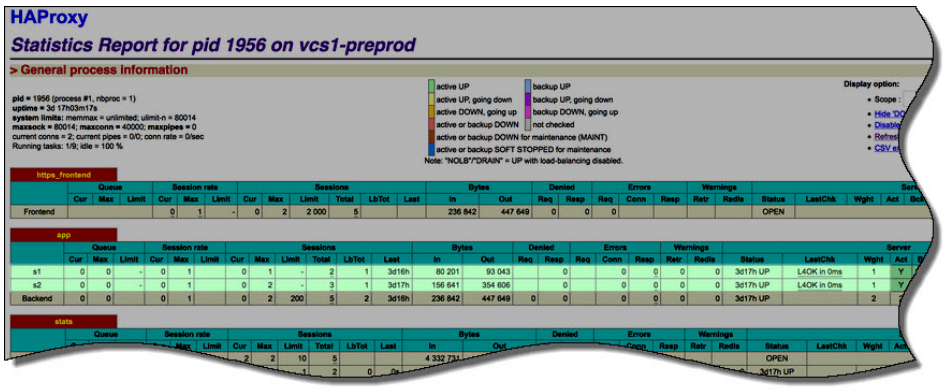

18 Repeat Step 15 through 17 on the **Slave** server.

Install keepalived on Each VCS Console VM

1 On the **Master** server, open the **/etc/sysctl.conf** file.

```
[root@vcsconsole ~]# vi /etc/sysctl.conf
```

2 Locate the following two fields and set the value to **1**. This will enable IP forwarding.

Note: Add these fields if they are not defined in the file by default.

```
net.ipv4.ip_forward = 1
net.ipv4.ip_nonlocal_bind = 1
```

3 Save and close the file.

4 Run the following command to update the server with new configuration.

```
[root@vcsconsole ~]# sysctl -p
```

5 Enter the following command to install **keepalived**.

```
[root@vcsconsole ~]# yum install keepalived
```

6 When prompted to confirm the installation, type **y**.

7 Repeat Steps 1 through 6 on the **Slave** server.

8 On the **Master** server, copy the keepalived configuration file to a backup file.

```
[root@vcsconsole ~]# mv /etc/keepalived/keepalived.conf
/etc/keepalived/keepalived.conf.def
```

9 Enter the following command to create anew keepalived.conf file in a text editor.

```
[root@vcsconsole ~]#  vi /etc/keepalived/keepalived.conf
```

10 Copy and paste the following content into the file.

```
! Configuration File for keepalived
#MASTER
#10.90.47.21
global_defs {
   lvs_id LoadBalancer01
}

vrrp_instance VI_1 {
    state MASTER
    interface eth0
    virtual_router_id 51
    priority 100
    advert_int 1
    authentication {
        auth_type PASS
        auth_pass 1111
    }
```

```
        virtual_ipaddress {
           10.90.47.191
        }
}

virtual_server 10.90.47.191 443 {
    delay_loop 6
    lb_algo rr
    lb_kind NAT
    nat_mask 255.255.252.0
    persistence_timeout 50
    protocol TCP

    real_server 10.90.47.21 443 {
        weight 1
        SSL_GET {
            url {
              path /
              digest d41d8cd98f00b204e9800998ecf8427e
            }
            connect_timeout 10
            nb_get_retry 50
            delay_before_retry 3
        }
    }
}
```

11 Update the following fields:

Note: The field that you need to edit are also shown in bold above.

Field	Value
state	MASTER/SLAVE **Note**: Enter the appropriate value for the server you are configuring.
virtual_route_id	A unique ID for a VCS Console cluster **Note**: This ID must be unique to each VCS Console cluster. Use the same ID for the MASTER and the SLAVE.
virtual_ipaddress	VIP address
virtual_server	VIP address

Field	Value
nat_mask	Netmask for VIP address
real_server	IP address for eth0 for the MASTER server

12 Save and close the file.

13 Repeat Steps 8 through 12 on the **Slave** server.

Important: Make sure to set the state field to **SLAVE**. All other fields are the same as for the MASTER server.

14 On the **Master** server, enter the following command to start the keepalived service.

```
[root@vcsconsole ~]# service keepalived start
```

15 Repeat Step 14 on the **Slave** server.

16 Enter the following command on the **Master** server to verity that the eth0 interface is configured.

Note: Ensure that the VIP is only plumbed on the Master server.

```
[root@vcsconsole ~]# ip add sh eth0
```

17 Repeat Step 16 on the **Slave** server.

18 Configure keepalived to start from boot up on the **Master** server.

```
[root@vcsconsole ~]$ sudo chkconfig keepalived on
```

19 Repeat Step 18 on the **Slave** server.

20 Enter the following command on the **Master** server to verify that the node took the expected role. The output should say "Entering MASTER STATE".

```
[root@vcsconsole ~]# grep Keepalived_vrrp /var/log/messages
```

Example:

```
Jun 13 12:01:40 vcs1-prod Keepalived_vrrp[10376]: VRRP_Instance(VI_1) Received lower prio advert, fo
rcing new election
Jun 13 12:01:41 vcs1-prod Keepalived_vrrp[10376]: VRRP_Instance(VI_1) Entering MASTER STATE
Jun 13 12:01:41 vcs1-prod Keepalived_vrrp[10376]: VRRP_Instance(VI_1) setting protocol VIPs.
Jun 13 12:01:41 vcs1-prod Keepalived_vrrp[10376]: VRRP_Instance(VI_1) Sending gratuitous ARPs on eth
0 for 172.20.35.20
Primary node:     "Entering MASTER STATE"
Secondary node: "Entering BACKUP STATE"
```

21 Repeat Step 20 on the Slave server. The output should say "Entering SLAVE STATE".

22 From a Web browser, enter the following command to test the VCS Console HA.

URL Syntax: https://[VIP]

Notes:

- You do not need to append the port to this URL.
- You may need to click **Add Exception** multiple times before the VCS Console Web UI displays.

23 Go to ***Install the ECS VM*** (on page 79).

Install the ECS VM

The ECS node is the server where the solution services are installed. Complete the following procedures twice to install two ECS nodes. Before beginning these procedures, ensure that your system meets the following requirements.

Important: If any of these requirements have not been completed, please do so now.

- Oracle RAC is installed
- ECS database users successfully created
- Linux platform template built
- Admin Node deployed and running
- Consul nodes deployed and running
- VCS Consoles deployed and running
- Two IP addresses
  - ECS 1 IP
  - ECS 2 IP

Creating the ECS VM

Important: If you are using vSphere client to deploy virtual machines, you cannot create the ECS VM using a template. Refer to ***Procedures When Using vSphere Client*** (on page 221).

Complete the following procedures to deploy the ECS VM from the vSphere Web UI.

1. From the vCenter Web UI, click **VMs and Templates**.
2. Locate and select the CSCOlxplat template that was built using the ***Admin Node Installation Guide***.
3. Right-click the template and select **Deploy VM from this Template**. The Deploy From Template window opens.
4. In the text box, enter a name for the VM you are creating and then select the datacenter where it will be deployed. Click **Next.**
5. Select the appropriate ESXI host where you want to deploy the VM and click **Next**.
6. From the **Select virtual disk format** dropdown menu, maintain the **Same format as source** default. Then ensure that the appropriate datastore is selected.
7. Click **Next** and then click **Next** again.
8. Review the settings and click **Finish**.

Reconfiguring the ECS 3.0 VM

Important: If you are using vSphere client to deploy and configure VMs, refer to *Reconfiguring the Virtual Network Using vSphere Client* (on page 223).

Complete the following procedure to reconfigure the virtual hardware on the ECS node.

1. Select and right-click the ECS VM and then select **Edit Settings**. The Edit Settings window appears.
2. From the **CPU** text box, change the value to **4**.
3. From the **Memory** text box, change the value to **4** and then click the dropdown box next to the value and select **GB**.
4. From the **Hard disk 1** text box, change the value to **32 GB**.
5. Is the Oracle RAC set up on a different network?
   - If **yes**, go to step 6.
   - If **no**, go to Step 10.
6. From the **New device** dropdown menu, select **Network** and then click **Add**.
7. Click the arrow next to the **New Network** entry to view the details.

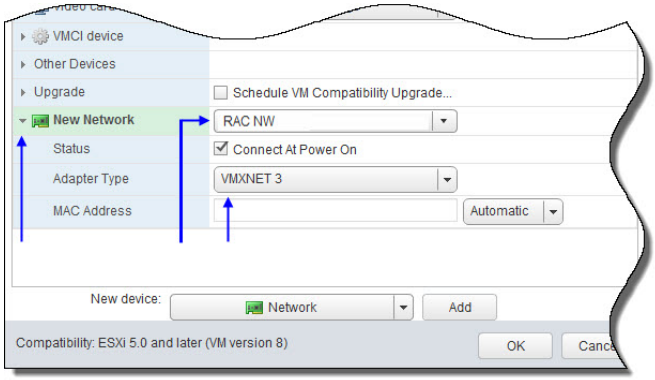

8. From the **New Network** dropdown menu, select the network associated with the RAC database.
9. From the **Adapter Type** dropdown menu, select **VMXNET 3**.
10. Click **OK**. The VM is reconfigured.

Configuring the ECS Network Interface With a Static IP

1. Select and right-click the ECS VM and select **Power On**.
2. Select and right-click the ECS VM again and select **Open Console**. A VMware console window opens in a new tab.
3. Log into VM with the following credentials.

   Username: admin

   Password: [password created for Linux platform template]
4. Enter the following command to change to **root** user.

   ```
   [admin@platform ~]$ sudo -i
   ```
5. Type the following command line utility to configure the network and the DNS settings. The Select Action window appears.

   ```
   [root@platform ~]# system-config-network
   ```
6. Select **Device configuration** and press **Enter**.
7. Highlight **eth0** and press **Enter**.
8. Press the **Tab** key until you **Use DHCP i**s highlighted. Then press the **Spacebar** to unselect this option.
9. Tab to each field to enter the following values.

   Note: DNS entries are optional.

   - **Static IP**
   - **Netmask**
   - **Default gateway IP**
   - **Primary DNS Server**
   - **Secondary DNS Server**
10. Verify that **Peer DNS** is selected.
11. Press the **Tab** key until you **Controlled by NetworkManager i**s highlighted. Then press the **Spacebar** to unselect this option.
12. Press **Tab** to highlight **Ok** and press **Enter**. The Select A Device window appears.
13. Click **Save**. The Select Action window appears.
14. Click **Save&Quit**.
15. Restart the network service to start the interface.

    ```
    [root@platform ~]# service network restart
    ```
16. Is the RAC in a different subnet than the VM?
    - If **no**, you have completed this procedure.
    - If **yes**, repeat Steps 4 through 14 to configure a secondary interface for eth1.

17 Verify that you can ping the RAC database.

18 Did you successfully ping the RAC database?

- If **yes**, you have completed this procedure.
- If **no**, troubleshoot your network or call Cisco Services.

Deploying the ECS VM

1 Using an SSH client, log into the ECS VM.

2 On the ECS VM, type the following command to create a **/var/tmp/staging** directory.

```
[root@ecs ~]# mkdir /var/tmp/staging
```

3 Change to the **/var/tmp/staging** directory.

```
[root@ecs ~]# cd /var/tmp/staging
```

4 Copy the **cisco-vcs-deployment** zip file from the Admin node to this directory on the ECS server.

Note: Substitute the IP address of your Admin node for the [Admin_IP] entry.

Command Syntax:

```
scp -Crp admin@[Admin_IP]:/opt/cisco/software/admin_node/
cisco-vcs-deployment-*.zip .
```

Example:

```
[root@ecs staging]# scp -Crp
admin@10.90.44.70:/opt/cisco/software/admin_node/
cisco-vcs-deployment-*.zip .
```

5 Unzip the deploy file and then change to the **scripts** directory.

```
[root@ecs staging]# unzip cisco-vcs-deployment-*.zip
[root@ecs staging]# cd cisco-vcs-deployment-*/scripts
```

6 Modify the ECS environment file, **ecs.envfile**. A description of each field in the file is shown below.

Field	Value
app_package_list	Packages to install
install_labcerts	Set to false
persistence_port	RAC database listening port; set to 1535
persistence_type	Database engine name; set to oracle
persistence_name	Database name; set to CABHE
service_auth_method	Set to twoway
service_protocol	Set to https
service_startup_time	Maintain default value of 1200
consul_datacenter	Consul datacenter name

Field	Value
admin_node	Admin Node name or IP address
persistence_host	RAC scan IP address
persistence_user	RAC ECS database user name
persistence_password	RAC ECS database user password
consul_servers	All Consul server IP addresses separated by a comma (no spaces)
consul_encrypt	Consul encryption key
hostname	ECS VM hostname

Default ecs.envfile

```
app_package_list=RegisterService,RegisterService_UI,CpeManagement,CpeManagement_UI,Loa
dPIMSService,OAMService,OAMService_UI,RPSService,RPSService_UI,ReportService_UI,SNMPTr
apConfigSvc,SNMPTrapSvc,SNMPTrapConfigUI
install_labcerts=false
persistence_port=1535
persistence_type=oracle
persistence_name=CABHE
service_startup_time=1200
consul_datacenter=dc1
admin_node=
persistence_host=
persistence_user=
persistence_password=
consul_servers=
consul_encrypt=
```

Example: ecs.envfile

```
app_package_list=RegisterService,RegisterService_UI,CpeManagement,CpeManagement_UI,Loa
dPIMSService,OAMService,OAMService_UI,RPSService,RPSService_UI,ReportService_UI,SNMPTr
apConfigSvc,SNMPTrapSvc,SNMPTrapConfigUI
install_labcerts=false
persistence_port=1535
persistence_type=oracle
persistence_name=CABHE
service_auth_method=twoway
service_protocol=https
service_startup_time=1200
consul_datacenter=dc1
admin_node=10.90.44.70
persistence_host=172,20.36.5
persistence_user=VCSPROD_ECS
persistence_password=vcsprod
consul_servers=10.90.47.33,10.90.47.34,10.90.47.35
consul_encrypt=414V9ZJWsKg/RrGSV+1+qw==
hostname=ecs-prod
```

7 Save and close the file.

8 Enter the following command to execute the ECS deploy script.

```
[root@ecs staging]# ./deploy-ecs.sh --envfile=ecs.envfile 2>&1
| sudo tee /var/log/deploy-ecs.log
```

Results: When the script completes, the ECS VM reboots.

Configuring X.509 Certificates for TLS Encryption on the ECS VM

Important: The certificate and key pair should have been generated when deploying and configuring the Admin Node. If they have not yet been created, go to the following chapters in the ***Admin Node Installation Guide*** to create them now.

- Chapter 5: Create Environment Files for NextX Nodes
- Chapter 6: Create NextX X.509 Root CA Certificates

This section includes the procedure to transfer the certificate/key pair generated and saved on the Admin Node to the respective node. Once transferred, a procedure is included to verify the configuration of the certificate.

Transferring X.509 Certificates From the Admin Node to the ECS Node

Important: This procedure must be executed for certificates that were generated from an internal root CA or from an external CA.

Complete the following steps to transfer the certificate files created for the ECS nodes from the Admin Node to each respective ECS node.

Note: This procedure is executed on the Admin Node.

1. From a terminal window, log into the Admin Node as **admin** user and then enter the following command to change to **root** user.

   ```
   [admin@adminnode ~]$ sudo -i
   ```

2. Enter the following command to change to the **/opt/cisco/ca** directory.

   ```
   [root@adminnode ~]# cd /opt/cisco/ca
   ```

3. Enter the following command and press **Enter** to transfer the appropriate certificate and key pair to the ECS node.

   Command Syntax:

   ```
   [root@adminnode ca]# ./manageCerts -P [absolute_path_to_cert]
   [absolute_path_to_key] [ECS_IP]
   ```

   Example:

   ```
   [root@adminnode ca]# ./manageCerts -P
   /etc/pki/CA/certs/ecs.domain.pem
   /etc/pki/CA/private/ecs.domain.key 10.90.47.246
   ```

   Notes:

   - Replace [cert] with the location of the node certificate file (e.g. /etc/pki/CA/certs/[CA.pem]) on the Admin Node.
   - Replace [key] with the location of the node private key file (e.g. /etc/pki/CA/private/[CA.key]) on the Admin Node.

- Replace [IP] with the IP address of the ECS node, which is the IP address defined as IP.1 in the [hostname].env file on the Admin Node.

4 Were you prompted to verify the SSH RSA key fingerprint:

- If **yes**, type **yes** and press **Enter**. Then go to the next step.
- If **no**, refer to the results section of the next step.

5 When prompted, enter and then re-enter the admin password for the ECS node.

Results:

- The certificate associated private key and truststore are distributed to the ECS node.
- Two-way authentication configuration is applied to the ECS.
- The ECS jboss service is restarted.
- A **./manageCerts finished** message displays.

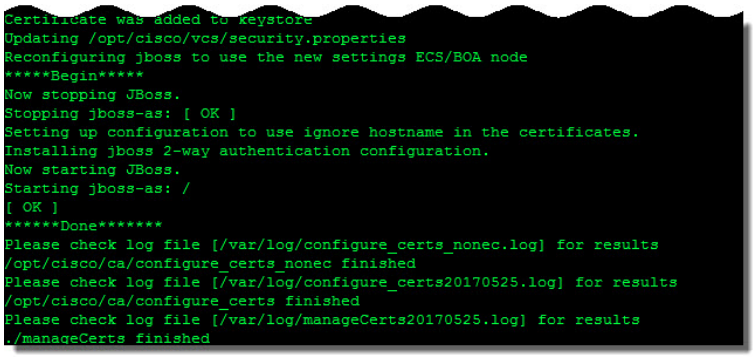

6 Review the logs in the **/var/log** directory.

7 Go to the next section.

Verifying the ECS Certificate Configuration

Once the ECS node is configured, you can execute the **checkConfig** script to verify the configuration. Complete the following steps to verify these parameters.

Note: You should still be logged into the Admin Node as the **admin** user in the /opt/cisco/ca directory.

1 As **admin** user on the Admin Node, enter the following command to check the certificate configuration for the ECS node in which certificates have been generated. A validation of the certificate files occurs.

Command Syntax:

```
sudo ./checkConfig -s [hostname].env
```

Example:

```
[admin@adminnode ca]$ sudo ./checkConfig -s ecs.env
```

```
[root@adminnode ca]# ./checkConfig -s ecs.env
==========20170531.121830==========
./checkConfig -s ecs.env

===Checking node at IP 10.90.47.159 from ecs.env
Validating files in check/10.90.47.159
Checking ecs at 10.90.47.159
10.90.47.159 has consul enabled
10.90.47.159 has consul running
10.90.47.159 has jboss-as enabled
10.90.47.159 has jboss-as running
Checking consul config etc/consul/config.json
Checking CA cert etc/pki/tls/cacert.pem
Checking server cert etc/pki/tls/certs/server.crt
Checking server key etc/pki/tls/private/server.key
Checking security properties opt/cisco/vcs/security.properties
Checking JBOSS config opt/jboss-as/standalone/configuration/standalone-vcs.xml
openssl verify -CAfile /opt/cisco/ca/check/10.90.47.159/etc/pki/tls/cacert.pem /opt/cisco/ca/che
ck/10.90.47.159/etc/pki/tls/certs/server.crt
/opt/cisco/ca/check/10.90.47.159/etc/pki/tls/certs/server.crt: OK

Please check log file [/var/log/checkConfig20170531.log] for results

./checkConfig finished
```

2 Did any errors display?

- If **yes**, review the **/var/log/checkConfig[date].log** file to remedy the issue. Then repeat Step 1. When the issues are corrected, you have completed this procedure.
- If **no**, go to the next section.

Configuring ECS Services

Complete the following procedure on the ECS as **root** user to configure ECS services.

1 Open the **/etc/jboss-as/conf.d/jboss-as-vcs.conf** file to define the bind address to the ECS IP address.

```
[root@ecs ~]# vi /etc/jboss-as/conf.d/jboss-as-vcs.conf
```

Example:

```
JAVA_OPTS="$JAVA_OPTS -Djboss.bind.address=10.90.47.97"
```

2 Enter the following command to configure JBoss.

```
[root@ecs ~]# /opt/cisco/jboss-config/setup.sh --auth two-way
```

Results:

- This script updates the /opt/jboss-as/standalone/configuration/standalone-vcs.xml file.
- JBoss is started.

3 Were you returned to a root prompt?

- If **yes**, go to the next step.
- If **no**, press **Ctrl+C** to return to the root prompt. Then go to the next step.

4 Enter the following command to enable the JBoss startup script.

```
[root@ecs ~]# chkconfig jboss-as on
```

5 Enter the following command to generate the ECS Web UI layout on the VCS Console.

```
[root@ecs ~]#
/opt/cisco/ecs/installed/ECSCommon/generateLayout.sh create
```

6 Verify that the ECS Web UI layout was successfully created for all of the Web UIs.

Note: Ignore the **Error: EASService_UI does not install yet** message as EAS is not yet installed.

```
[root@ecs ~]# sudo grep -i "UI layout"
/opt/jboss-as/standalone/log/ECSUIutil.log
```

```
2017-03-29 08:46:50,106 INFO  RegisterService_UI layout has been created successfully
2017-03-29 08:46:58,146 INFO  OAMService_UI layout has been created successfully
2017-03-29 08:47:06,333 INFO  CpeManagement_UI layout has been created successfully
2017-03-29 08:47:14,033 INFO  RPSService_UI layout has been created successfully
```

7 Enter the following command to verify that the ECS host is now a member of the consul group.

```
[root@ecs ~]# consul members
```

Example:

```
Node                          Address             Status  Type    Build  Protocol  DC
vodwater                      10.90.46.232        alive   client  0.7.2  2         dc1
consul.domain.consul          10.90.47.108:8301   alive   server  0.7.2  2         dc1
consul.domain.consul2         10.90.47.109:8301   alive   server  0.7.2  2         dc1
consul.domain.consul3         10.90.47.177:8301   alive   server  0.7.2  2         dc1
ecs.domain.ecs1               10.90.47.159:8301   alive   client  0.7.2  2         dc1
vcsconsole.domain.vcsconsole1 10.90.47.145:8301   alive   client  0.7.2  2         dc1
vcsconsole.domain.vcsconsole2 10.90.47.189:8301   alive   client  0.7.2  2         dc1
```

8 Login to VCS Console Web UI and verify that the ECS services are listed in the VCS Console home page.

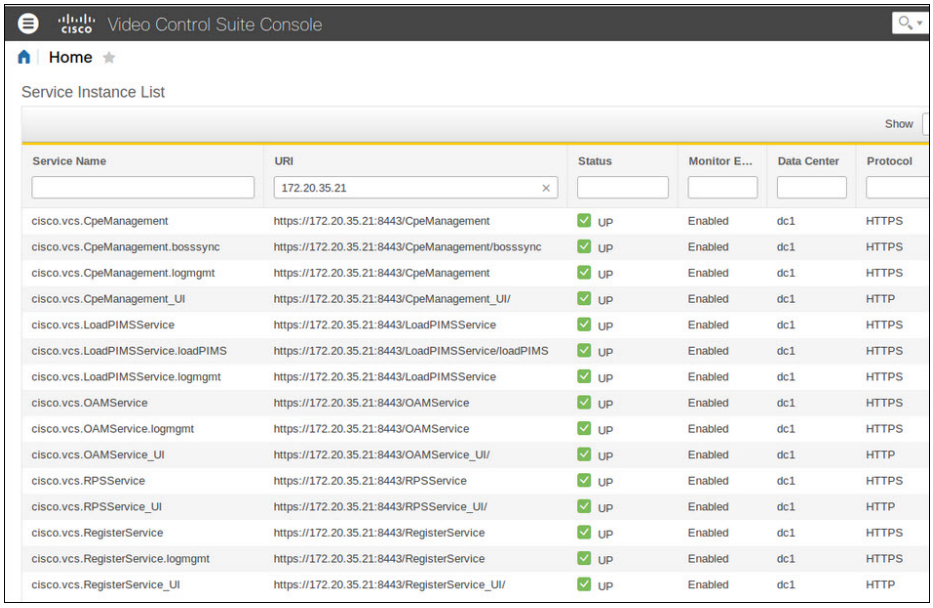

9 From the VCS Console, click **Control Plane > SNMP Trap Configuration UI**. The SNMP Trap Configuration window appears.

10 Click the **Add** button. The SNMP Trap Configuration window opens.

11 Enter the appropriate values for the following fields:

Field	Value
IP Address	VCS VIP address
Port Number	162
SNMP Version	v2c
Community	Public

Example:

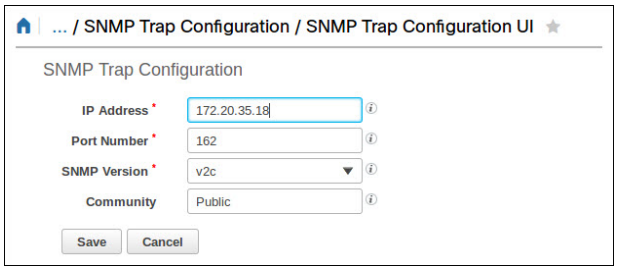

12 Click **Save**. If successful, the following message appears in the lower, right corner of the Web UI window.

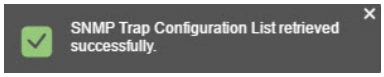

13 Reload the window to see the new SNMP configuration in the SNMP Trap Configuration List.

14 Click **Console Admin > SSO Configuration** to configure single sign on (SSO). The Configure SSO Details window appears.

15 In the Server Host text box, enter the VIP for the VCS Console.

Example:

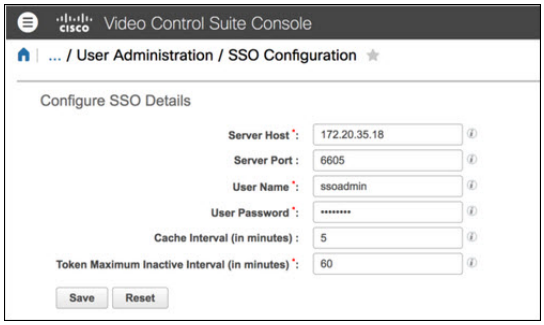

16 Click **Save**.

Configuring snmpd on the ECS Node

Complete the following procedure to configure snmpd on the ECS node to monitor the following services:

- Consul
- ECS - JBoss

1 Enter the following command to install the **ServiceManager** package on your ECS node.

```
[root@ecs ~]# yum install ServicesManager
```

2 When prompted to confirm the installation, type **y** and press **Enter**. When the installation is finished, a **Complete!** message displays.

3 As **root** user, enter the following command to edit the **/etc/snmp/snmpd.conf** file in a text editor.

```
[root@ecs ~]# vi /etc/snmp/snmpd.conf
```

4 Add the following lines to the end of the file:

```
# Monitor consul process and send traps
view    systemview    included   .1.3.6.1.4.1.1429
view    systemview    included   .1.3.6.1.4.1.2021

rocommunity public
rwuser admin
iquerySecName admin
agentSecName admin
proc consul 1 1
proc ecssvcsmon 1 1

monitor -r 5 -u admin -i -o .1.3.6.1.4.1.2021.2.1.2.1 -o
.1.3.6.1.4.1.2021.2.1.101.1 "Consul process has stopped." -o
.1.3.6.1.4.1.2021.2.1.100.1 .1.3.6.1.4.1.2021.2.1.100.1 != 0
monitor -r 5 -u admin -i -o .1.3.6.1.4.1.2021.2.1.2.1 -o
.1.3.6.1.4.1.2021.2.1.101.1 "Consul process is running." -o
.1.3.6.1.4.1.2021.2.1.100.1 .1.3.6.1.4.1.2021.2.1.100.1 == 0

monitor -r 5 -u admin -i -o .1.3.6.1.4.1.2021.2.1.2.2 -o
.1.3.6.1.4.1.2021.2.1.101.2 "ECS processes have stopped." -o
.1.3.6.1.4.1.2021.2.1.100.2 .1.3.6.1.4.1.2021.2.1.100.2 != 0
monitor -r 5 -u admin -i -o .1.3.6.1.4.1.2021.2.1.2.2 -o
.1.3.6.1.4.1.2021.2.1.101.2 "ECS processes is running." -o
.1.3.6.1.4.1.2021.2.1.100.2 .1.3.6.1.4.1.2021.2.1.100.2 == 0
```

```
trapsess -v 2c -c public <<VCS Console HA IP>>:162
```

5 Save and close the file.

6 Enter the following command to restart the **snmpd** service.

```
[root@ecs ~]# service snmpd restart
```

7 Enter the following command to verify that the **snmpd** is running.

```
[root@ecs ~]# service snmpd status
```

Output:

```
snmpd (pid  4609) is running...
```

8 Enter the following command to enable service monitoring.

```
[root@ecs ~]# /opt/cisco/servicesmgr/bin/enableMonitoring.sh
```

9 Enter the following command to verify that the monitoring processes are running.

```
[root@ecs ~]# ps -ef | grep -i mon
```

```
dbus      1369     1  0 Jun08 ?        00:00:00 dbus-daemon --system
root      2247     1  0 Jun08 ?        00:15:17 monit
root      5891     1  0 Jun10 ?        00:00:00 /usr/sbin/console-kit-daemon --no-daemon
root     12628 12584  0 16:07 pts/0    00:00:00 grep -i mon
```

10 Have you installed and configured two ECS nodes?

- If **no**, go to ***Creating the ECS VM*** (on page 79) and complete all of the procedures, up to and including this section, to install and configure the second ECS node.
- If **yes**, go to ***Create the BOA VM*** (on page 92).

Configuring the DHCT De-Register Option

The CpeManagement.properties file on the ECS node allows you to configure how CPE devices are de-registered from the system. The setting, deleteOnDeregister, includes the following two options.

- **false**: (Default option)
  - Deletes the ECID *associated* with the CPE MAC addresses of the CC/DHCT from the CPEMS database.
  - Does ***NOT*** delete the CPE MAC addresses of the CC/DHCT from the CPEMS database.
  - Deletes the CPE MAC addresses of the CC/DHCT from the associated EC database.

  Result: The CPE MAC addresses remain in the cloud repo but are removed from the targeted EC. A second delete transaction would be required to remove the CPE from the CPEMS database (cloud repo).

- **true:**
  - The ECID and the CPE MAC addresses are removed from the CPEMS database (cloud repo).
  - The CPE MAC addresses are removed from the associated EC database.

  Result: The CPE MAC addresses are not present in the CPEMS database or the EC database.

Complete the following steps to change the deleteOnDeregister value to true and remove the ECID and MAC address from both the CPEMS and the EC databases.

1. From the *primary* ECS node, login as **admin** user.
2. Enter the following command to set the **deleteOnDeregister** value to **true**.

   ```
   [admin@ecs ~]$ sudo vi /opt/
   jboss-as/standalone/configuration/CpeManagement.properties
   ```
3. Go to the **deleteOnDerigister** line and change the value to **true**.

   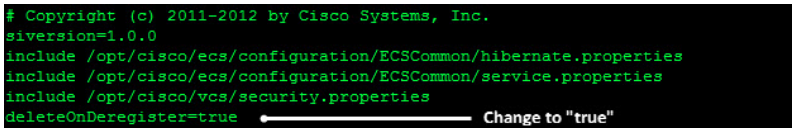

4. Save and close the file.
5. Enter the following command to restart **jboss**.

   ```
   [admin@ecs ~]$ sudo service jboss-as status
   ```
6. Repeat Steps 1 through 5 on the *secondary* ECS node.

Create the BOA VM

The Business Support System/Operational Support System (BSS/OSS) adapter (BOA) service allows you manage the billing system interface for each head-end control system in one central location.

Complete the following procedures to install two BOA nodes. Before beginning these procedures, ensure that your system meets the following requirements.

Important: If any of these requirements have not been completed, please do so now.

- Linux platform template built
- Admin Node deployed and running
- Consul nodes deployed and running
- VCS Consoles deployed and running
- Two IP addresses are required for failover purposes
  - BOA 1 IP
  - BOA 2 IP

Creating the BOA VM

Important: If you are using vSphere client to deploy virtual machines, you cannot create the BOA VM using a template. Refer to ***Procedures When Using vSphere Client*** (on page 221).

Complete the following procedures to deploy the BOA VM from the vSphere Web UI.

1. From the vCenter Web UI, click **VMs and Templates**.
2. Locate and select the CSCOlxplat template that was built using the ***Admin Node Installation Guide***.
3. Right-click the template and select **Deploy VM from this Template**. The Deploy From Template window opens.
4. In the text box, enter a name for the VM you are creating and then select the appropriate datacenter where it will be deployed. Then click **Next.**
5. Select the appropriate ESXi host where you want to deploy the VM and click **Next**.
6. From the **Select virtual disk format** dropdown menu, maintain the **Same format as source** default. Then ensure that the appropriate datastore is selected.
7. Click **Next** and then click **Next** again.
8. Review the settings and click **Finish**.

Reconfiguring the BOA VM

Complete the following procedure to reconfigure the virtual hardware on the BOA node.

1. Select and right-click the BOA VM and then select **Edit Settings**. The Edit Settings window appears.
2. From the **CPU** text box, change the value to **4**.
3. From the **Memory** text box, change the value to **4** and then click the dropdown box next to the value and select **GB**.
4. From the **Hard disk 1** text box, change the value to **32 GB**.
5. Will billing be set up on a different network?
   - If **yes**, go to step 6.
   - If **no**, go to Step 10.
6. From the **New device** dropdown menu, select **Network** and then click **Add**.
7. Click the arrow next to the **New Network** entry to view the details.
8. From the **New Network** dropdown menu, select the network associated with the billing network.
9. From the **Adapter Type** dropdown menu, select **VMXNET 3**.
10. Click **OK**. The VM is reconfigured.

Configuring the BOA Network Interface With a Static IP

1. Select and right-click on the BOA VM and select **Power On**.
2. Select and right-click the BOA VM again and select **Open Console**. A VMware console window opens in a new tab.
3. Log into VM with the following credentials.

   Username: admin

   Password: [password created for Linux platform template]
4. Change to **root** user.

   ```
   [admin@platform ~]$ sudo -i
   ```
5. Type the following command line utility to configure the network and the DNS settings. The Select Action window appears.

   ```
   [root@platform ~]# system-config-network
   ```
6. Select **Device configuration** and press **Enter**.
7. Highlight **eth0** and press **Enter**.
8. Press the **Tab** key until you **Use DHCP i**s highlighted. Then press the **Spacebar** to unselect this option.

9 Tab to each field to enter the following values.

Note: DNS entries are optional.

- **Static IP**
- **Netmask**
- **Default gateway IP**
- **Primary DNS Server**
- **Secondary DNS Server**

10 Verify that **Peer DNS** is selected.

11 Press the **Tab** key until you **Controlled by NetworkManager i**s highlighted. Then press the **Spacebar** to unselect this option.

12 Press **Tab** to highlight **Ok** and press **Enter**. The Select A Device window appears.

13 Is billing on a different network than eth0?

- If **yes**, go to the next step to configure the eth1 network interface.
- If **no**, go to Step 23.

14 Press the arrow key on your keyboard until **<New Device>** is highlighted and then press **Enter**. The Network Configuration window displays.

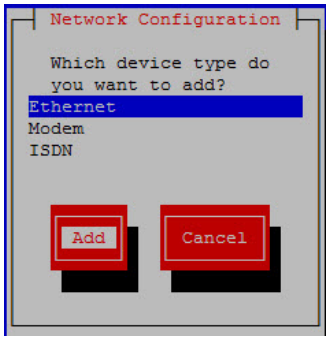

15 With **Ethernet** highlighted, press the **Tab** key until **Add** is highlighted and then press **Enter**.

16 In the **Name** text box, type **eth1** and press **Tab**.

17 In the **Device** text box, type **eth1** and press **Tab** twice.

18 Verify that **Use DHCP** is *not* selected.

19 Tab to each of the following fields to enter the appropriate values for your system.

- **Static IP**
- **Netmask**
- **Default gateway IP**
- **Primary DNS Server**
- **Secondary DNS Server**

20 Verify that **Peer DNS** is selected.

21 Press the **Tab** key until you **Controlled by NetworkManager i**s highlighted. Then press the **Spacebar** to unselect this option.

22 Press **Tab** to highlight **Ok** and press **Enter**. The Select A Device window appears.

23 Press **Tab** until **Save** is highlighted and then press Enter. The Select Action window appears.

24 Click **Save&Quit**.

25 Restart the network service to start the interface.

```
[root@platform ~]# service network restart
```

Deploying the BOA VM

1 Using an SSH client, log into the BOA VM and change to **root** user.

2 On the BOA VM, type the following command to create a **/var/tmp/staging** directory.

```
[root@boa ~]# mkdir /var/tmp/staging
```

3 Change to the **/var/tmp/staging** directory.

```
[root@boa ~]#  cd /var/tmp/staging
```

4 Copy the **cisco-vcs-deployment** zip file from the Admin node to this directory on the BOA VM.

Note: Substitute the IP address of your Admin node for the [Admin_IP] entry.

Command Syntax:

```
scp -Crp admin@[Admin_IP]:/opt/cisco/software/admin_node/
cisco-vcs-deployment-*.zip .
```

Example:

```
[root@boa ~]#  scp -Crp
admin@10.90.44.70:/opt/cisco/software/admin_node/
cisco-vcs-deployment-*.zip .
```

5 Unzip the deploy file and then change to the **scripts** directory.

```
[admin@boa staging]$ unzip cisco-vcs-deployment-*.zip
[admin@boa staging]$ cd cisco-vcs-deployment-*/scripts
```

6 Modify the BOA environment file, **boa.envfile**. A description for each field in the file is shown below.

Notes: Append **rpcbind,vsftpd** to the app_package_list if BOA will use RPC BOSS. Inclusion of these package names will result in the installation of the rpcbind and vsftpd package on the BOA node.

Field	Value
app_package	List of packages
install_labcerts	Set to false

Field	Value
service_startup_time	Maintain default value of 1200
consul_datacenter	Consul datacenter name
admin_node	Admin Node name or IP address
consul_servers	All Consul server IP addresses separated by a comma (no spaces)
consul_encrypt	Consul encryption key
hostname	BOA hostname

Default boa.envfile

```
app_package_list=billingAdaptor,BillingAdaptorUI
install_labcerts=false
service_startup_time=1200
consul_datacenter=dc1
admin_node=
consul_servers=
consul_encrypt=
```

Example: boa.envfile

```
app_package_list=billingAdaptor,BillingAdaptorUI
install_labcerts=false
service_startup_time=1200
consul_datacenter=dc1
admin_node=10.90.44.70
consul_servers=10.90.47.33,10.90.47.34,10.90.47.35
consul_encrypt=414V9ZJWsKg/RrGSV+1+gw==
hostname=boa
```

7 Save and close the file.

8 Enter the following command to execute the BOA deploy script.

```
[root@boa ~]# ./deploy-boa.sh --envfile=boa.envfile 2>&1 |
sudo tee /var/log/deploy-boa.log
```

Results: When the script completes, the BOA VM reboots.

9 Log back into the BOA node as **admin** user.

Enabling the Transfer of Pay-Per-View Reports

Important: Complete this procedure only if BOA will use RPC BOSS. If BOA is not using RPC BOSS, skip this procedure and go to the next section.

To transfer pay-per-view (PPV) reports, an FTP user must be created and vsftpd must be configured. Complete this procedure to enable this feature.

1 If you are not root user, enter the following command to change to **root** user.

```
[admin@boa ~]$ sudo -i
```

2 Enter the following command to create the FTP user.

```
[root@boa ~]# useradd -m dncsftp
```

3 Open the **/etc/vsftpd/vsftp.conf** file in a text editor.

```
[root@boa ~]# vi /etc/vsftpd/vsftpd.conf
```

4 Uncomment the following line to resemble the following output.

```
chroot_local_user=YES
```

5 Verify that the following entry is present.

```
userlist_deny=NO
```

6 Save and close the file.

Configuring BOA

1 Create the following file using a text editor.

```
[root@boa ~]# vi /etc/jboss-as/conf.d/boa.conf
```

2 Insert the following content into the **/etc/jboss-as/conf.d/boa.conf** file.

Note: For the Djboss.bind.address entry, replace [IP Address] with the actual IP address of the node.

```
JBOSS_CONFIG=standalone-vcs.xml
JAVA_OPTS="-Xms2g -Xmx4g -XX:MaxPermSize=2g"
JAVA_OPTS="$JAVA_OPTS -d64"
JAVA_OPTS="$JAVA_OPTS -Djava.net.preferIPv4Stack=true"
JAVA_OPTS="$JAVA_OPTS -Djava.awt.headless=true"
JAVA_OPTS="$JAVA_OPTS -Djboss.bind.address=[IP Address]"
export JAVA_OPTS
```

3 Save and close the file.

4 Enter the following command to change the ownership of the file to **jboss:jboss**.

```
[root@boa ~]# chown jboss:jboss /etc/jboss-as/conf.d/boa.conf
```

Important: If any other jboss configuration file (i.e. *.conf) file exists besides boa.conf, either remove it or rename it so it will not be used.

5 Type the following command to setup the BOA config.properties file. A > prompt appears.

```
[root@boa ~]# bash -c 'cat >
/opt/cisco/billingadaptor/conf/config.properties' <<EOF
```

6 Copy and paste the following content into the file.

Note: The keystore and truststore parameters will be updated by the manageCerts utility during the ***Transferring X.509 Certificates From the Admin Node to the BOA Node*** (on page 99) procedure.

```
connections.consul.ip=localhost
connections.consul.port=8500
connections.sync.soap.longtimeoutinseconds=30
sec.keystore.keystorePath=
sec.keystore.password=
sec.keystore.truststorePath=
sec.twoWayAuthEnabled=true
```

```
serviceDirectory.type=consul
serviceDirectory.datacenter=dc1
serviceDirectory.ip=localhost
serviceDirectory.port=8500
serviceDirectory.requiredFlag=true
service.name.CpeManagement.bosssync=cisco.vcs.CpeManagement.bosssync
service.name.LoadPIMSService=cisco.vcs.LoadPIMSService.loadPIMS
EOF
```

7 Change the ownership of the **config.properties** file to **jboss:jboss**.

```
[root@boa ~]# chown jboss:jboss /opt/cisco/billingadaptor/conf/config.properties
```

8 Open the **/etc/jboss-as/conf.d/jboss-as-vcs.conf** file in a text editor and change the **boss.bind.address** value to the actual IP address of the node.

Example:

```
JAVA_OPTS="$JAVA_OPTS -Djboss.bind.address=10.90.47.180"
```

9 Save and close the **/etc/jboss-as/conf.d/jboss-as-vcs.conf** file.

10 Enter the following command to start the JBOSS application.

```
[root@boa ~]# service jboss-as start
```

11 Execute the following command to configure the JBoss application to start automatically during system boot.

```
[root@boa ~]# chkconfig jboss-as on
```

12 Execute the following command to verify that the JBoss application is configured to automatically start during system boot.

```
[root@boa ~]# chkconfig jboss-as --list
```

Note: Init levels 2 through 5 should be "on" for the jboss-as application.

Example:

```
jboss-as 0:off  1:off   2:on    3:on    4:on    5:on    6:off
```

13 Enter the following command to deploy billingAdaptor.

```
[root@boa ~]# /opt/cisco/billingadaptor/bin/billingadaptor.sh deploy
```

14 As **root** user, enter the following command to check the current status of BOA.

```
[root@boa ~]# /opt/cisco/billingadaptor/bin/billingadaptor.sh status
```

Successful BOA Deploy Output:

```
[standalone@localhost:9999 /]     deployment-info --name=BillingAdaptor.war
NAME              RUNTIME-NAME       PERSISTENT ENABLED STATUS
BillingAdaptor.war BillingAdaptor.war true       true    OK
[standalone@localhost:9999 /]     exit
```

Unsuccessful BOA Deploy Output:

```
[standalone@localhost:9999 /]     deployment-info --name=BillingAdaptor.war
JBAS014807: Management resource '[("deployment" => "BillingAdaptor.war")]' not found
```

15 Was the deployment successful?

- If **yes**, you have completed this procedure. Go to the next section.
- If **no**, go to the next step.

16 Complete the following steps to undeploy and then redeploy BOA.

a Type the following command to undeploy BOA.

```
[root@boa ~]#
/opt/cisco/billingadaptor/bin/billingadaptor.sh undeploy
```

b Type the following command to redeploy BOA.

```
[root@boa ~]#
/opt/cisco/billingadaptor/bin/billingadaptor.sh deploy
```

c Repeat Step 14 to check the status of the deployment.

Configuring X.509 Certificates for TLS Encryption on the BOA VM

Important: The certificate and key pair should have been generated when deploying and configuring the Admin Node. If they have not yet been created, go to the following chapters in the ***Admin Node Installation Guide*** to create them now.

- Chapter 5: Create Environment Files for NextX Nodes
- Chapter 6: Create NextX X.509 Root CA Certificates

This section includes the procedure to transfer the certificate/key pair generated and saved on the Admin Node to the respective node. Once transferred, a procedure is included to verify the configuration of the certificate.

Transferring X.509 Certificates From the Admin Node to the BOA Node

Important: This procedure must be executed for certificates that were generated from an internal root CA or from an external CA.

Complete the following steps to transfer the certificate files created for the BOA nodes from the Admin Node to each respective BOA node.

Note: This procedure is executed on the Admin Node.

1 From a terminal window, log into the Admin Node as **admin** user and then enter the following command to change to **root** user.

```
[admin@adminnode ~]$ sudo -i
```

2 Enter the following command to change to the **/opt/cisco/ca** directory.

```
[root@adminnode ~]# cd /opt/cisco/ca
```

3 Enter the following command and press **Enter** to transfer the appropriate certificate and key pair to the BOA node.

Command Syntax:

```
[root@adminnode ca]# ./manageCerts -P [absolute_path_to_cert]
[absolute_path_to_key] [BOA_IP]
```

Example:

```
[root@adminnode ca]# ./manageCerts -P
/etc/pki/CA/certs/boa.domain.pem
/etc/pki/CA/private/boa.domain.key 10.90.47.246
```

Notes:

- Replace [cert] with the location of the node certificate file (e.g. /etc/pki/CA/certs/[CA.pem]) on the Admin Node.
- Replace [key] with the location of the node private key file (e.g. /etc/pki/CA/private/[CA.key]) on the Admin Node.
- Replace [IP] with the IP address of the BOA node, which is the IP address defined as IP.1 in the [hostname].env file on the Admin Node.

4 Were you prompted to verify the SSH RSA key fingerprint:

- If **yes**, type **yes** and press **Enter**. Then go to the next step.
- If **no**, refer to the results section of the next step.

5 When prompted, enter and then re-enter the admin password for the BOA node.

Results:

- The certificate associated private key and truststore are distributed to the BOA node.
- TLS is enabled for RPS encryption.
- The BOA service is restarted.
- A **./manageCerts finished** message displays.

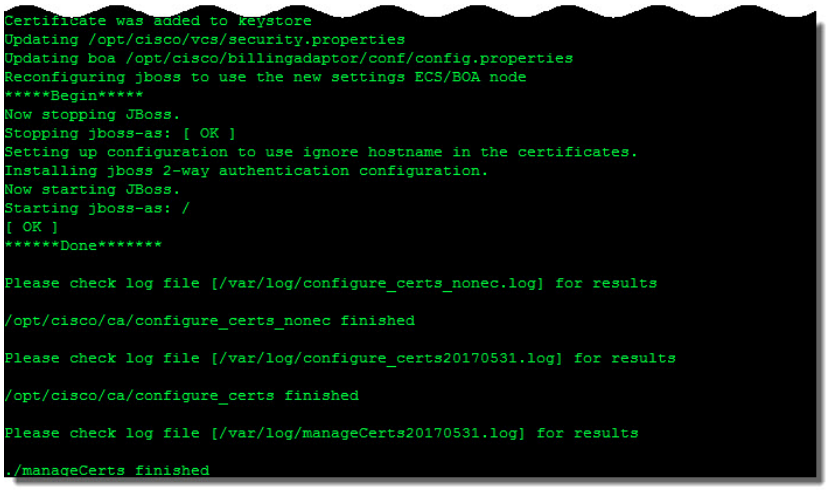

6 Review the logs in the **/var/log** directory.

7 Go to the next section.

Verifying the BOA Certificate Configuration

Once the BOA node is configured, you can execute the **checkConfig** script to verify the configuration. Complete the following steps to verify these parameters.

Note: You should still be logged into the Admin Node as the admin user in the /opt/cisco/ca directory.

1 As **admin** user on the Admin Node, enter the following command to check the certificate configuration for the BOA node in which certificates have been generated. A validation of the certificate files occurs.

Command Syntax:

```
sudo ./checkConfig -s [hostname].env
```

Example:

```
[admin@adminnode ca]# sudo ./checkConfig -s boa1.env
```

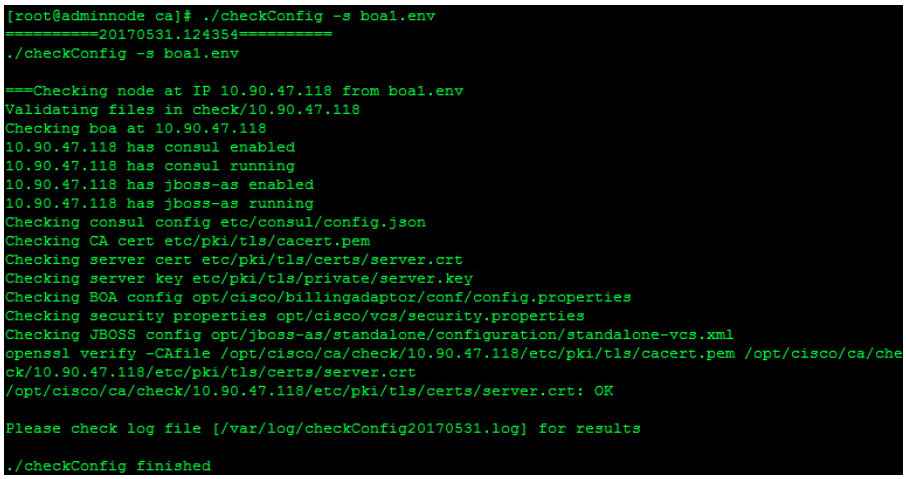

2 Did any errors display?

- If **yes**, review the **/var/log/checkConfig[date].log** file to remedy the issue. Then repeat Step 1. When the issues are corrected, you have completed this procedure.
- If **no**, go to the next section.

Configuring the BOA Web UI

Complete the following steps as **root** user on the BOA node to configure the BOA Web UI to communicate over https.

1 Open the **consoleasservice.properties** in a text editor.

2 `[root@boa ~]# vi /opt/cisco/installed/BillingAdaptorUI-*/BillingAdaptorUI/WEB-INF/classes/consoleasservice.properties`

3 Go to the **##UI Plug In Configurations Console as Service** section and enter the following values for the fields shown below.

```
CAS_SERVICE_NAME=BillingAdaptorUI
VCS_PORT=8443
VCS_PROTOCOL=https
REGISTER_SERVICE_SLEEP=60000
```

4 Save and close the file.

5 Enter the following command to deploy the BillingAdaptorUI war file.

```
[root@boa ~]# touch /opt/jboss-as/standalone/deployments/BillingAdaptorUI.war.deployed
```

6 For further configuration, refer to the ***Installing the Billing Adaptor (BOA) User's Guide***.

Configuring snmpd on the BOA Node

Complete the following procedure to configure snmpd on the BOA node to monitor the following services:

- Consul
- BOA - JBoss

1 Enter the following command to install the **ServicesManager** package on your BOA node.

```
[root@boa ~]# yum install ServicesManager
```

2 When prompted to confirm the installation, type **y** and press **Enter**. When the installation is finished, a **Complete!** message displays.

3 As **root** user, enter the following command to edit the **/etc/snmp/snmpd.conf** file in a text editor.

```
[root@boa ~]# vi /etc/snmp/snmpd.conf
```

4 Add the following lines to the end of the file:

```
# Monitor consul process and send traps
view    systemview    included   .1.3.6.1.4.1.1429
view    systemview    included   .1.3.6.1.4.1.2021

rocommunity public
rwuser admin
iquerySecName admin
agentSecName admin
proc consul 1 1
proc ecssvcsmon 1 1

monitor -r 5 -u admin -i -o .1.3.6.1.4.1.2021.2.1.2.1 -o
.1.3.6.1.4.1.2021.2.1.101.1 "Consul process has stopped." -o
.1.3.6.1.4.1.2021.2.1.100.1 .1.3.6.1.4.1.2021.2.1.100.1 != 0
monitor -r 5 -u admin -i -o .1.3.6.1.4.1.2021.2.1.2.1 -o
.1.3.6.1.4.1.2021.2.1.101.1 "Consul process is running." -o
.1.3.6.1.4.1.2021.2.1.100.1 .1.3.6.1.4.1.2021.2.1.100.1 == 0

monitor -r 5 -u admin -i -o .1.3.6.1.4.1.2021.2.1.2.2 -o
.1.3.6.1.4.1.2021.2.1.101.2 "BOA processes have stopped." -o
.1.3.6.1.4.1.2021.2.1.100.2 .1.3.6.1.4.1.2021.2.1.100.2 != 0
monitor -r 5 -u admin -i -o .1.3.6.1.4.1.2021.2.1.2.2 -o
.1.3.6.1.4.1.2021.2.1.101.2 "BOA processes is running." -o
.1.3.6.1.4.1.2021.2.1.100.2 .1.3.6.1.4.1.2021.2.1.100.2 == 0

trapsess -v 2c -c public <<VCS Console HA IP>>:162
```

5 Save and close the file.

6 Enter the following command to restart the **snmpd** service.

```
[root@boa ~]# service snmpd restart
```

7 Enter the following command to verify that the **snmpd** is running.

```
[root@boa ~]# service snmpd status
```

Output:

```
snmpd (pid  4609) is running...
```

8 Enter the following command to enable service monitoring.

```
[root@boa ~]# /opt/cisco/servicesmgr/bin/enableMonitoring.sh
```

9 Enter the following command to verify that the monitoring processes are running.

```
[root@boa ~]# ps -ef | grep -i mon
```

```
dbus      1369     1  0 Jun08 ?        00:00:00 dbus-daemon --system
root      2247     1  0 Jun08 ?        00:15:17 monit
root      5891     1  0 Jun10 ?        00:00:00 /usr/sbin/console-kit-daemon --no-daemon
root     12628 12584  0 16:07 pts/0    00:00:00 grep -i mon
```

10 Have you installed and configured two BOA nodes?

- If **yes**, go to the next step.
- If **no**, go to ***Creating the BOA VM*** (on page 92) and complete all of the procedures, up to and including this section, to install and configure the second BOA node.

11 Is this a new ECS 3.0 installation?

- If **yes**, you have completed the installation. Go to ***Verifying ECS Functionality*** (on page 140).
- If **no** and this is an ECS migration, go to ***Migrate to ECS*** (on page 105).

5

Migrate to ECS

Important: You must have completed all of the procedures in the ***Install and Configure the ECS System*** (on page 43) chapter prior to executing the procedures in this chapter.

This chapter provides the procedures for migrating ECS 2.0 to the ECS system. The procedures must be completed in the order in which they are written.

The unload and load scripts are located on the SR 3.0 ECS services node with the exception of AlarmService and Reports Server. The scripts for these two services are on the VCS Console node.

You will be required to copy the unload scripts for Oracle to the Oracle server node. After executing the unload scripts, you will transfer the unloaded data to the destination system where the data will be loaded onto the Oracle or MySQL database for your SR 3.0 system.

The following list of services are supported:

- RegisterService
- OAMService
- ReportDiscoveryService
- LoadPIMS
- RPSService
- AlarmService
- Reports Service

In This Chapter

Migrating Data From the CMC to the VCS Console

This section provides the procedures for migrating CMC data to the new VCS Console node. This migration consists of the following two steps.

1. Data is extracted from the CMC and stored in a json file using the data export script.
2. Data is loaded into the VCS Console using an import script that includes the json file.

Data That is Migrated to the VCS Console

USERS

- Users that have at least one associated User Group present in the VCS Console database.
- Users are not migrated if it already exists in the VCS Console database (e.g. root/admin).

AAAMode

- AAAMode configuration data

RADIUS/Tacacs/LDAP

- RADIUS configuration data

  Note: Local Interface IP is changed at server start up.
- Tacacs configuration data

  Note: Local Interface IP is changed at server start up.
- LDAP configuration including certificate data

Authorization Configuration

- User Group Attributes
- User Groups Mapping

SFTP User

- SFTP User Account

Login Disclaimer

- Login Disclaimer Text

Audit Log

- Audit Log Settings

Logging Configuration

- Message Levels
- Logging Configuration (AAA, Admin, Database, GUI and System)
- Log File Settings

Network Element Access

- Network Element Access configuration data

User - User Group Association

- User and User Group Association are only migrated for User Groups present in VCS Console.

Role Privilege

- Any privileges added to a Role
- Any privileges removed from a Role

Data That is Not Migrated to the VCS Console

USERGROUPS

- UserGroups

  Note: UserGroups are automatically created by plugin installs (e.g. ECS UI install)

User - User Group Association

- Group associations (i.e. Message-infra-Mgr, Cloud-DVR-Management

Password Policy

- Password policy data (due to new password policies in ECS 3.0)

Prerequisites

Important: CMC data migration to the VCS Console is only supported on Oracle Database Servers.

- VCS Console 4.0-16 or later
- ECS Services and Web UIs are registered with VCS Console
- A Search filter must be configured if an LDAP server is present on the CMC

Verifying Network Connectivity for the CMC Data Export

The export utility is designed to be executed from the VCS Console node assuming there is network connectivity to the CMC. Complete the following procedure to verify network connectivity between the CMC and the VCS Console.

1 Obtain the following information from the ECS 2.0 system.

Note: In an HA environment, either VCS Console Node can be used.

- ECS 2.0 CMC HA IP (Same as the CMC Web UI IP): _____
- ECS 3.0 VCS Console IP: _____

2 Enter the following commands to verify if there is network connectivity between the VCS Console and the CMC.

```
[root@rac2-mgmt1 ~]# ssh admin@<VCS Console IP>
[root@vcsconsole ~]# ssh root@<CMC HA IP>
```

Example: Shows that network connectivity is successful using CMC HA IP = 172.20.35.41 and VCS Console IP = 72.20.10.5

```
$ ssh admin@172.20.35.41
admin@172.20.35.41's password:
Last login: Tue Mar 21 09:51:08 2017 from 172.16.16.100
$ ssh root@172.20.10.5 date
root@172.20.10.5's password:
Tue Mar 21 11:13:44 EDT 2017
```

3 Does network connectivity exist between the CMC and the VCS Console?

- If **yes**, go to *Exporting CMC Data With Network Connectivity to the VCS Console* (on page 109).
- If **no**, go to *Exporting CMC Data Without Network Connectivity to the VCS Console* (on page 111).

Exporting CMC Data With Network Connectivity to the VCS Console

1 As **root** user on the VCS Console, change to the **/opt/data/vcsdatahandler/conf** directory.

```
[root@vcsconsole ~]# cd  /opt/data/vcsdatahandler/conf
```

2 Edit the **exportConfig.properties** file in a text editor and update the following fields with database values specific to your system.

Note: This information resides in the CMC database on the ECS 2.0 system.

```
[root@vcsconsole conf]# vi exportConfig.properties
```

Field	Value
dbUsername	wcsdba
dbPassword	wcs123
dbHost	CMC HA IP address (VIP)
dbPort	1522

Field	Value
dbDriver	Oracle
dataFile	Json file where data will be exported

Default exportConfig.properties File

```
dbUsername=
dbPassword=
dbHost=
dbPort=
dbSid=
dbDriver=oracle.jdbc.driver.OracleDriver
```

Example exportConfig.properties File

```
dbUsername=wcsdba
dbPassword=wcs123
dbHost=172.20.10.5
dbPort=1522
dbSid=wcs
dbDriver=oracle.jdbc.driver.OracleDriver
dataFile=/tmp/cmc.json
```

3 Save and close the file.

4 Enter the following command to execute the **exportCMCData.sh** script.

Note: This script will extract the data from the CMC and store it in the json file specified in the exportConfig.properties file.

```
[root@vcsconsole conf]# ../scripts/exportCMCData.sh
```

```
Writing logs to File /opt/data/vcsdatahandler/logs/migration/migration.log

Export Started - Wed Feb 08 16:51:52 EST 2017
#######################READING FROM DATABASE#######################

#################### DB Details ################################
DATABASE Host : 172.20.10.5
DATABASE Port : 1522
DATABASE UserName : wcsdba
DATABASE Password : ********************
DATABASE Service : wcs
DATABASE Driver : oracle.jdbc.driver.OracleDriver
DATABASE Trying to get connection from Server..
DATABASE Connection Successful..
################################################################
READING USERS....
READING USER GROUP ASSOCIATION....
READING USER ROLE ASSOCIATION....
READING NETWORK ELEMENT ACCESS ASSOCIATION....
READING ROLE ATTRIBUTES MAPPING WITH EXTERNAL SERVERS....
READING AAAMODE....
READING ROLE ASSOCIATION WITH EXTERNAL SERVERS....
READING AUTHENTICATION MANAGERS, AUDIT and LOGGING SETTINGS....
READING FTP AND DISCLAIMER DETAILS....
READING AAA LDAP SERVER FILE DETAILS....
READING ROLE PRIVILEGE ASSOCIATION....
###################READING FROM DATABASE COMPLETED################
WRITING TO FILE -------> /tmp/cmc.json


Data Export Completed Successfully
Export Completed - Wed Feb 08 16:51:53 EST 2017
```

5 Monitor the output until the script completes. A **Data Export Completed Successfully** message appears.

6 Open the **/opt/data/vcsdatahandler/logs/migration/migration.log** to review the export of the data.

```
[root@vcsconsole conf]# less ../logs/migration/migration.log
```

7 Go to *Importing the CMC Data Into the VCS Console* (on page 113).

Exporting CMC Data Without Network Connectivity to the VCS Console

1 On the VCS Console, enter the following command to compress the **vcsdatahandler** directory.

```
[root@vcsconsole conf]# cd /opt/data
[root@vcsconsole data]# tar -hcf /tmp/vcsdatahandler.tar
vcsdatahandler
```

2 Copy the **/tmp/vcsdatahandler.tar** file from the VCS Console to the **/tmp** directory on the CMC node.

Important: Because there is not direct connectivity between the VCS Console and the CMC, you cannot scp the file. You will need to find another method to transfer the file.

3 Using an SSH client, log into with the CMC with the HA IP address as **root** user.

Command Syntax:

```
ssh root@[CMC HA IP]
```

4 Create the **/opt/data** directory on the CMC node and then extract the **vcsdatahandler.tar** file into it.

```
[root@rac2-mgmt1 data]# mkdir /opt/data
[root@rac2-mgmt1 data]# cd /opt/data
[root@rac2-mgmt1 data]# tar xvf /tmp/vcsdatahandler.tar
```

5 On the VCS Console, add the following JAVA_EXEC entry into the **exportCMCData.sh** script so it will function properly with the CMC node.

JAVA_EXEC="/usr/java/jdk1.6.0.51/bin/java"

Important: Add the JAVA_EXEC variable exactly as shown and in the exact line where it is shown in this example.

```
[root@vcsconsole data]# vi
/opt/data/vcsdatahandler/scripts/exportCMCData.sh
```

```
#!/bin/bash
DATAHANDLER_HOME=/opt/data/vcsdatahandler
DATAHANDLER_LOG_DIR="$DATAHANDLER_HOME/logs"
CLASSPATH="${DATAHANDLER_HOME}/conf:${DATAHANDLER_HOME}/lib/*:"
DATAHANDLER_CONF_DIR="$DATAHANDLER_HOME/conf"

JAVA_EXEC="java"

if [ -e /opt/web/vcsconsole/bin/java_vcsconsole/bin/java ]
  then
       JAVA_EXEC="/opt/web/vcsconsole/bin/java_vcsconsole/bin/java"
fi
JAVA_EXEC="/usr/java/jdk1.6.0_51/bin/java"   <-- Add this entry
$JAVA_EXEC -cp $CLASSPATH \
                       -Djava.security.egd=file:/dev/./urandom \
                       -DlogFile=${DATAHANDLER_LOG_DIR}/migration/migration.lo
g \
               com.cisco.vcs.migration.exp.ExportDataMain --config ${DATAHANDL
ER_CONF_DIR}/exportConfig.properties
```

6 Update the CMC database information in the **exportConfig.properties** file.

```
[root@vcsconsole data]# vi
/opt/data/vcsdatahandler/conf/exportConfig.properties
```

Field	Value
dbUsername	wcsdba
dbPassword	wcs123
dbHost	localhost
dbPort	1522
dbSid	wcs
dbDriver	Oracle
dataFile	Json file where data will be exported

Default exportConfig.properties File

```
dbUsername=
dbPassword=
dbHost=
dbPort=
dbSid=
dbDriver=oracle.jdbc.driver.OracleDriver
dataFile=
```

Example exportConfig.properties File

```
dbUsername=wcsdba
dbPassword=wcs123
dbHost=localhost
dbPort=1522
dbSid=wcs
dbDriver=oracle.jdbc.driver.OracleDriver
dataFile=/tmp/cmc.json
```

7 Enter the following command to execute the **exportCMCData.sh** script.

Note: This script extracts the data from the CMC and stores it in the json file specified in the exportConfig.properties file.

```
[root@vcsconsole data]#
/opt/data/vcsdatahandler/scripts/exportCMCData.sh
```

8 Monitor the output until the script completes. A **Data Export Completed Successfully** message appears.

```
DATABASE Connection S
##################################################################
READING USERS....
READING USER GROUP ASSOCIATION....
READING USER ROLE ASSOCIATION....
READING NETWORK ELEMENT ACCESS ASSOCIATION....
READING ROLE ATTRIBUTES MAPPING WITH EXTERNAL SERVERS....
READING AAAMODE....
READING ROLE ASSOCIATION WITH EXTERNAL SERVERS....
READING AUTHENTICATION MANAGERS, AUDIT and LOGGING SETTINGS....
READING FTP AND DISCLAIMER DETAILS....
READING AAA LDAP SERVER FILE DETAILS....
READING ROLE PRIVILEGE ASSOCIATION....
###################READING FROM DATABASE COMPLETED#################
WRITING TO FILE -------> /tmp/cmc.json


Data Export Completed Successfully
Export Completed - Wed Feb 08 16:51:53 EST 2017
```

9 Open the **/opt/data/vcsdatahandler/logs/migration/migration.log** to review the export of the data.

```
[root@vcsconsole data]# less
/opt/data/vcsdatahandler/logs/migration/migration.log
```

10 Copy the **/tmp/cmc.json** file from the CMC node to the **/tmp** directory on the VCS Console.

Important: Because there is not direct connectivity between VCS Console and the CMC, you cannot scp the file. You will need to find another method to transfer the file.

11 Go to the next section.

Importing the CMC Data Into the VCS Console

1 Login to both VCS Console servers as root user and enter the following command to stop the vcsconsole service.

```
[root@vcsconsole data]# service vcsconsole stop
```

2 `[root@vcsconsole2 data]# service vcsconsole stop`

3 Open the **/opt/data/vcsdatahandler/conf/importConfig.properties** file in a text editor and update the following fields to reflect the database information for your system.

```
[root@vcsconsole data]# vi
/opt/data/vcsdatahandler/conf/importConfig.properties
```

Field	Value
dbUsername	VCS Console database user
dbPassword	VCS Console database user password
dbHost	Oracle RAC Scan IP address
dbPort	1535
dbSid	CABHE
dbDriver	Oracle
dataFile	Json file where data will be exported; same file as where the CMC data was exported

Default importConfig.properties File

```
dbUsername=
dbPassword=
dbHost=
dbPort=
dbSid=
dbDriver=oracle.jdbc.driver.OracleDriver
dataFile=
```

Example importConfig.properties File

```
dbUsername=ECS2MIGR_VCSCONSOLE
dbPassword=ecs2migr
dbHost=172.20.36.5
dbPort=1535
dbSid=CABHE
dbDriver=oracle.jdbc.driver.OracleDriver
dataFile=/tmp/cmc.json
```

4 Enter the following command to execute the **importCMCData.sh** script. This script will load the CMC data into the VCS Console.

```
[[root@vcsconsole data]#
/opt/data/vcsdatahandler/scripts/importCMCData.sh
```

5 When prompted to confirm the import, type **y**. The migration of the data begins.

6 Monitor the output until the script completes. A **Data Import Completed Successfully** message appears.

Note: Some errors may appear for users that are already in the system or for groups that were associated with groups that do not exist in the VCS Console. You can ignore these messages.

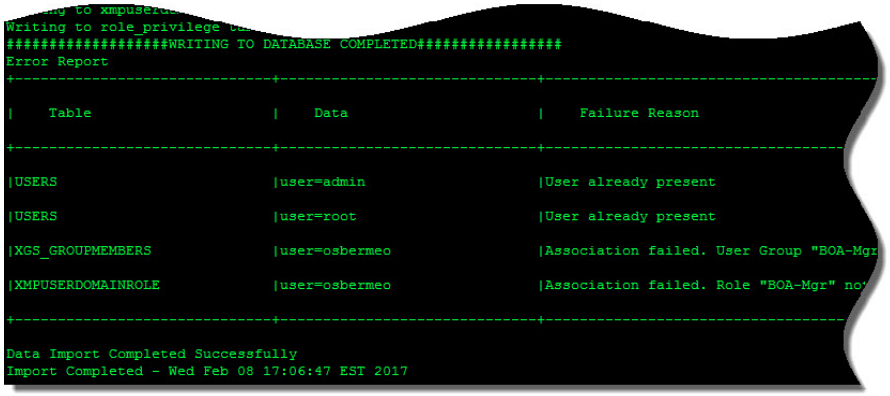

7 Run the following command on each VCS Console nodes to start the vcsconsole service.

```
[root@vcsconsole data]# service vcsconsole start
```

Migrating Alarm Settings

The following scripts reside in the /opt/AlarmManager/scripts on the VCS Console. They will be used to migrate the alarm settings to the ECS system.

- unloadoraecsdb_alarm.sh
- loadunloadscripts.readme
- loadoraecsdb_alarm.sh

Complete the following procedure to migrate the alarm settings.

1 On the VCS Console, change to the **/opt/AlarmManager/scripts** directory.

```
[root@vcsconsole data]# cd /opt/AlarmManager/scripts
```

2 Using scp, copy the **unloadecsdb_alarm.sh** script from the VCS Console to the **/tmp** directory on the ECS 2.0 CMC server (i.e. MGMT node) with lumos database.

Command Syntax:

```
scp unloadecsdb_alarm.sh root@[ECS_2.0_IP]:/tmp
```

Example:

```
[root@vcsconsole scripts]# scp unloadecsdb_alarm.sh
root@10.90.181.249:/tmp
```

3 On the ECS 2.0 CMC node, enter the following command to access the lumos database.

```
[root@rac2-mgmt1 tmp]# su - oracle
```

4 Enter the following command to verify that you are in the **/common/oracle** directory.

```
[oracle@rac2-mgmt1 ~]$ pwd
```

5 Are you in the **/common/oracle** directory?

- If **yes**, go to the next step.
- If **no**, type `cd /common/oracle` to change to that directory.

6 Enter the following commands, one at a time, to source the environment.

Note: Press **Enter** after typing each command.

```
[oracle@rac2-mgmt1 ~]$ setenv DB_HOME /opt/oracle
[oracle@rac2-mgmt1 ~]$ setenv ORACLE_BASE /opt/oracle/base
[oracle@rac2-mgmt1 ~]$ setenv ORACLE_HOME
/opt/oracle/base/product/11.2.0/dbhome_1
[oracle@rac2-mgmt1 ~]$ setenv ORACLE_SID wcs
[oracle@rac2-mgmt1 ~]$ setenv SQLPATH /opt/oracle
[oracle@rac2-mgmt1 ~]$ setenv PATH  ${ORACLE_HOME}/bin:${PATH}
[oracle@rac2-mgmt1 ~]$ setenv LD_LIBRARY_PATH
${ORACLE_HOME}/lib
```

7 Type **exit** to exit the lumos database.

8 Enter the following command on the ECS 2.0 CMC node to unload the alarms from the ECS database.

Command Syntax:

```
./unloadoraecsdb_alarm.sh wcsdba wcs123 [acct_username]
[acct_password] [IP] [port] [SID]
```

Example:

```
[root@rac2-mgmt1 tmp]# ./unloadoraecsdb_alarm.sh wcsdba wcs123
10.90.181.249 1522 wcs
```

Result: When the script completes, an **All tables were SUCCESSFULLY UNLOADED FOR AlarmSetting service** message appears, along with the location of the log file and the unload file.

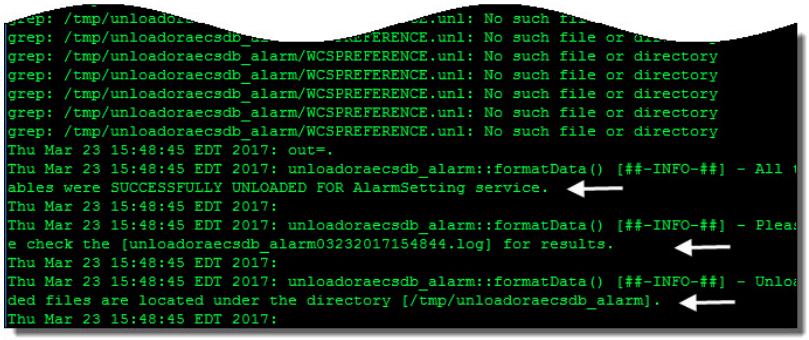

9 Go to the **/tmp/unloadoraecsdb_alarm/** directory.

```
[root@rac2-mgmt1 ~]# cd /tmp/unloadoraecsdb_alarm
```

10 Type `ls -ltr` and verify that the following files were generated from Step 8.

- USERPREFERENCECONFIG.unl
- WCSPREFERENCE.unl

11 Using scp, copy the **USERPREFERENCECONFIG.unl** to the /tmp/unloadoraecsdb_alarm/ directory on the ECS 3.0 Oracle RAC node.

12 Log into the ECS Oracle RAC 3.0 node and enter the following commands.

```
[root@nextxecs4rac4a ~]$ su - oracle
```

13 Enter the following command to source the environment variable.

```
[root@nextxecs4rac4a ~]$ . /opt/oracle/CABHE01.env
```

14 Change to the **/tmp/unloadoraecsdb_alarm** directory.

```
[root@nextxecs4rac4a ~]$ cd /tmp/unloadoraecsdb_alarm
```

15 Enter the following command to migrate the alarms to the ECS 3.0 Oracle database.

Command Syntax:

```
./loadoraecsdb_alarm.sh "acct un" "pw" "ip" 'port" "SID"
```

Example:

```
[oracle@nextxecs4rac4a unloadoraecsdb_alarm]$
./loadoraecsdb_alarm.sh migrate_alarms migrate 204.3.65.49
1535 CABHE01
```

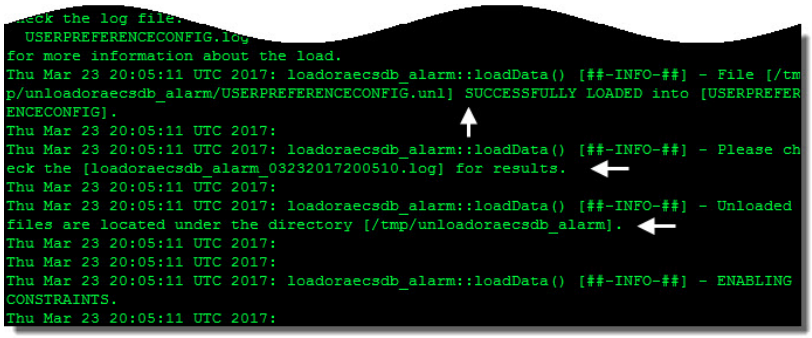

Result: When the script completes, a **SUCCESSFULLY LOADED** message appears, along with the location of the log file and the unload file.

Migrating the CPEMS Batch File

The cpeBatchImport.sh script resides in the on the ECS 3.0 node and is used to migrate the CPEMs batch files from ECS 2.0 to the ECS system. The batch tar files are located in the /common/adminsftp/upgrade/cpeManagement directory on the ECS 2.0 node.

Once the batch files are imported to ECS 3.0 they will be saved and stored in the database.

Note: The batch files are not saved to the filesystem.

1. Copy the batch tar files from the ECS 2.0 MGMT node to the **/opt/cisco/ecs/installed/CpeManagement/scripts** directory on the ECS 3.0 node.
2. On the ECS 3.0 node, change to the **/opt/cisco/ecs/installed/CpeManagement/scripts** directory.

   ```
   [root@migrate-ecs ~]# cd
   /opt/cisco/ecs/installed/CpeManagement/scripts
   ```
3. On the ECS 3.0 node, execute the **cpeBatchImport.sh** script.

   Command Syntax:

   ```
   ./cpeBatchImport.sh <batch file path>
   ```

   Example:

   ```
   [root@migrate-ecs scripts]# ./cpeBatchImport.sh SP00876867.tar
   ```
4. When the script completes, review the **/tmp/cpeBatchImport.out** log file.

   ```
   [root@migrate-ecs scripts]# less /tmp/cpeBatchImport.out
   ```
5. Enter the following command to remove the batch tar file from the /opt/cisco/ecs/installed/CpeManagement/scripts since they are now loaded in the database.

   Command Syntax:

   ```
   rm /opt/cisco/ecs/installed/CpeManagement/scripts
   [filename.tar]
   ```

   Example:

   ```
   [root@migrate-ecs scripts]# rm
   /opt/cisco/ecs/installed/CpeManagement/scripts SP00876867.tar
   ```

Migrating Alarms Data

The following scripts reside in the /opt/AlarmManager/scripts directory on the VCS Console. They will be used to migrate the alarm settings to the ECS system.

- unloadoraecsdb_alarmdata.sh
- loadoraecsdb_alarmdata.sh

Complete the following procedure to migrate the alarm data.

1 On the VCS Console node, stop the **AlarmManagerService** service.

```
[root@vcsconsole ~]# /opt/AlarmManager/decapcontrol.sh stop
```

2 Enter the following command to verify the status of tomcat.

```
[root@vcsconsole ~]# ps -ef | grep -i tomcat
```

3 Is tomcat running?

- If **no**, go to Step 4.
- If **yes**, go to Step 5.

4 Enter the following command to kill the tomcat process.

Note: Substitute the PID number from the output in Step 3 for [tomcat_PID].

```
[root@vcsconsole ~]# kill -9 [tomcat_PID]
```

5 Change to the **/opt/AlarmManager/scripts** directory.

```
[root@vcsconsole ~]# cd /opt/AlarmManager/scripts
```

6 Using scp, copy the **unloadecsdb_alarmdata.sh** script from the VCS Console to the **/tmp** directory on the ECS 2.0 CMC server (i.e. MGMT node) with lumos database.

Command Syntax:

```
scp unloadecsdb_alarmdata.sh root@[ECS_2.0_IP]:/tmp
```

Example:

```
[root@vcsconsole ~]# sudo scp unloadecsdb_alarmdata.sh
root@10.90.181.249:/tmp
```

7 On the ECS 2.0 CMC node, enter the following command to access the lumos database.

```
[root@rac2-mgmt1 tmp]# su - oracle
```

8 Enter the following command to verify that you are in the **/common/oracle** directory.

```
[oracle@rac2-mgmt1 ~]$ pwd
```

9 Are you in the **/common/oracle** directory?

- If **yes**, go to the next step.
- If **no**, type `cd /common/oracle` to change to that directory.

10 Enter the following commands, one at a time, to source the environment.

Note: Press **Enter** after typing each source environment variable.

```
[oracle@rac2-mgmt1 ~]$ setenv DB_HOME  /opt/oracle
[oracle@rac2-mgmt1 ~]$ setenv ORACLE_BASE  /opt/oracle/base
[oracle@rac2-mgmt1 ~]$ setenv ORACLE_HOME
/opt/oracle/base/product/11.2.0/dbhome_1
[oracle@rac2-mgmt1 ~]$ setenv ORACLE_SID  wcs
[oracle@rac2-mgmt1 ~]$ setenv SQLPATH  /opt/oracle
[oracle@rac2-mgmt1 ~]$ setenv PATH  ${ORACLE_HOME}/bin:${PATH}
[oracle@rac2-mgmt1 ~]$ setenv LD_LIBRARY_PATH
${ORACLE_HOME}/lib
```

11 Type **exit** to exit the lumos database.

12 Enter the following command to execute the **unloadoraecsdb_alarmdata.sh** script on the ECS 2.0 CMC node. The script will unload the alarms from the ECS database.

Note: Do not specify the [dest_directory]. The unload file will be saved to the default directory which is /tmp/unloadoraecsdb_ecip.

Command Syntax:

```
./unloadoraecsdb_alarmdata.sh wcsdba wcs123 [mgmt_node_ip]
[port] [SID] [alarm_source ip] [new_alarm_source_ip]
[dest_directory]
```

Example:

```
[oracle@rac2-mgmt1 ~]$ sudo ./unloadoraecsdb_alarmdata.sh
wcsdba wcs123 10.90.181.249 1522 wcs 10.90.46.12 10.90.47.225
```

Result: When the script completes, an **All tables were SUCCESSFULLY UNLOADED FOR AlarmSettingData service** message appears, along with the location of the log file and the unload file.

13 Go to the **/tmp/unloadoraecsdb_ecip** directory.

```
[root@rac2-mgmt1 ~]# cd /tmp/unloadoraecsdb_ecip
```

14 Verify that the following two files were generated from Step 12.

- ALARM.unl
- EVENT.unl

15 Using SCP, copy the **ALARM.unl** and the **EVENT.unl** files to the **/tmp/unloadoraecsdb_alarmdata** on the ECS 3.0 Oracle RAC server.

16 From the VCS Console, use scp and copy the **loadoraecsdb_alarmdata.sh** to the **/home/oracle/scripts/alarm_migration_scripts** directory on the ECS 3.0 Oracle RAC server.

17 Log into the ECS Oracle RAC 3.0 node and enter the following command to change to **oracle** user.

```
[root@nextxecs4rac4a unloadoraecsdb_alarm]# su - oracle
```

18 Enter the following command to source the environment variable.

```
[oracle@nextxecs4rac4a ~]$ . /opt/oracle/CABHE01.env
```

19 Change to the **/home/oracle/scripts/alarm_migration_scripts** directory.

```
[oracle@nextxecs4rac4a ~]$ cd
/home/oracle/scripts/alarm_migration_scripts
```

20 Enter the following command to migrate the alarms data to the ECS 3.0 Oracle database.

Command Syntax:

```
./loadoraecsdb_alarmdata.sh [acct_username] [acct_password]
[IP] [port] [SID]
```

Example:

```
[oracle@nextxecs4rac4a alarm_migration_scripts]$
./loadoraecsdb_alarm.sh migrate_alarms migrate 204.3.65.49
1535 CABHE01
```

Result: When the script completes, a **SUCCESSFULLY LOADED** message appears, along with the location of the log file and the unload file.

Migrating RPS EC Device Jobs

The following scripts and README file reside in the /opt/cisco/ecs/installed/RPSService/scripts directory on the ECS 3.0 server. They will be used to migrate the RPS EC device jobs from ECS 2.0 to the ECS 3.0 system.

- unloadoraecsdb_rpss_byec.sh
- loadoraecsdb_rpss_byec.sh
- rpsFileImport.sh
- loadunloadscripts.readme

Complete the following procedure to migrate the RPS data.

1. On the ECS 3.0 server, use scp to copy the **unloadoraecsdb_rpss_byec.sh** script to the **/home/oracle/scripts/rps_migration_scripts** directory on the ECS 2.0 Oracle RAC server.
2. On the ECS 2.0 Oracle RAC server, change to **oracle** user.

   ```
   [root@nextxecs4rac4a ~]$ su - oracle
   ```
3. Enter the following command to source the environment.

   ```
   [oracle@legacyRAC ~]$ . /opt/oracle/CABHE01.env
   ```
4. Change to the **/home/oracle/scripts/rps_migration_scripts** directory.

   ```
   [oracle@legacyRAC ~]$ cd
   /home/oracle/scripts/rps_migration_scripts
   ```
5. Execute the following command to unload the RPS data to an unload file.

   Command Syntax:

   ```
   ./unloadoraecsdb_rpss_byec.sh [db_username] [db_password] [IP]
   [port] [SID] [ECID] [EC_7.0 IP] [EC_8,0 IP] [unload_directory]
   ```

   Example

   ```
   [oracle@legacyRAC rps_migration_scripts]$
   ./unloadoraecsdb_rpss_byec.sh RAC2ECSAPP RAC2ECSAPP
   10.90.181.243 1535 CABHE vodwater 10.90.46.12 10.90.47.225
   /home/oracle/scripts/rps_migration_scripts
   ```
6. Type the following command to list the contents of the directory.

   ```
   [oracle@legacyRAC rps_migration_scripts]$ ls -latr
   ```

   Result: The following files should be present.

   - RPOVERSIONTOELEMENTMAPPING.unl
   - PROVISIONINGELEMENT.unl
   - JOBSTATUS.unl
   - IMPORTSTATUSREPORT.unl
   - .ecid

7 Using scp, copy the **loadoraecsdb_rpss_byec.sh** script from the ECS 3.0 server to the **/home/oracle/scripts/rps_migration_scripts** directory on the ECS 3.0 Oracle RAC server.

8 Using scp, copy all of the **.unl** and **.ecid** files from the ECS 2.0 server to the **/home/oracle/scripts/rps_migration_scripts** directory on the ECS 3.0 Oracle RAC server.

9 On the ECS 3.0 Oracle RAC server, enter the following commands to change the ownership of the files copied from Step 8 to **oracle:dba**.

```
[root@nextxecs4rac4a ~] # cd
/home/oracle/scripts/rps_migration_scripts
[root@nextxecs4rac4a rps_migration_scripts]# chown oracle:dbs
*.*
```

10 Enter the following command to load the RPS data in the ECS 3.0 Oracle RAC server database.

Command Syntax:

```
./loadoraecsdb_rpss_byec.sh [db_username] db_password] [IP]
[port] [SID] [ECID [unload_directory]
```

Example:

```
[root@nextxecs4rac4a rps_migration_scripts]#
./loadoraecsdb_rpss_byec.sh migrate_ecs migrate 204.3.65.49
1535 CABHE vodwater /home/oracle/scripts/rps_migration_scripts
```

Migrating Reports

The following scripts reside in the /opt/cisco/reports/scripts directory on the ECS 3.0 server. They will be used to migrate reports data from ECS 2.0 to the ECS system.

- loadoraecsdb_reports.sh
- unloadoraecsdb_reports.sh

Complete the following procedure to migrate the reports data.

1 On the ECS 3.0 server, change to the **/opt/cisco/reports/scripts** directory.

```
[admin@migrate-ecs scripts]$ cd /opt/cisco/reports/scripts
```

2 Copy the **unloadoraecsdb_reports.sh** script to the SR 2.0 Oracle RAC server.

3 Log into the SR 2.0 Oracle RAC server.

Command Syntax:

```
ssh root@[Oracle RAC IP]
```

Example:

```
ssh root@204.3.102.11
```

4 Enter the following command to verify that the unload script is present.

```
[root@rac2-mgmt1 ~]# ls
```

5 Enter the following command to unload the reports for each service.

Command Syntax:

```
./unloadoraecsdb_reports.sh wcsdba wcs123 [mgmt_node_ip]
[port] [SID] [EC_hostname] [EC_IP] [EC_Web UI_URL]
```

Example:

```
[oracle@rac2-mgmt1 reports]$ ./unloadoraecsdb_reports.sh
wcsdba wcs123 10.90.181.249 1522 wcs vodwater 10.90.46.182
https://10.90.46.182
```

Results:

- A **All tables were SUCCESSFULLY UNLOADED FOR ReportSettings service** message displays.
- The following three files should be present in the /tmp/unloadoraecsdb_report directory: NE_DETAILS.unl, REPORT_SETTINGS.unl, SCHEDULE_REPORTS.unl

6 Copy the three files generated from Step 5 to the ECS 3.0 Oracle RAC server.

7 On the ECS 3.0 Oracle RAC server, execute the following command to load the reports data to the database.

Command Syntax:

```
./loadoraecsdb_reports.sh [ECS_DB_Username] [ECS_DB_password]
[Database_IP] [port] [Database_Name]
```

Example:

```
[oracle@C5openstackb report]$ ./loadoraecsdb_report.sh
s4condor_ecs s4condor_ecs 10.90.46.182 1535 CABHE01
```

Result: A Table REPORT_SETTINGS successful message displays.

8 Open the following file to review the log.

```
[oracle@C5openstackb report]$
/home/oracle/scripts/reports/REPORT_SETTINGS.log
```

Updating RADIUS, LDAP and RBAC Attributes

After the migration is complete, change the Radius Profile attributes from Conductor to the VCS Console. The attributes for each system are shown below. This change will also be required for LDAP if custom RBAC attributes are configured.

- ECS 3.0 VCS Console: VCS Console:role0=Root
- ECS 2.0 CMC Node: Conductor:role0=Root

In ECS 3.0, the purged log directory resides in a different location as shown below.

- ECS 3.0 VCS Console: /opt/vcs/logs/VCS Console
- ECS 2.0 CMC Node: /opt/CSCOcmc/conf/ifm/

Complete the following steps to update the directory path.

1 Login to the VCS Console.

2 Click the **Navigation** button, , and then click **Console Admin > Audit Setting**. The Audit Log Purge Settings window appears.

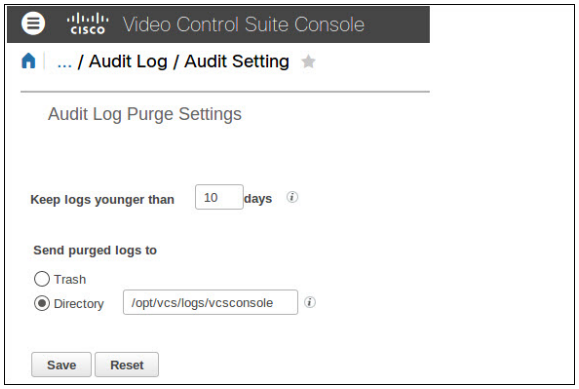

3 In the **Directory** text box, change this directory path **/opt/vcs/logs/vcsconsole**.

4 Click **Save**.

Expanding Storage on the Oracle RAC

Important: The Oracle RAC database ***does not*** need to be stopped. Both of the RAC VMs will need to be rebooted, but only one at a time. This procedure is more complicated to perform and results in a RAC disk configuration that is different from the standard installation from PowerShell script.

Once the ECS 2.0 system can be shutdown and deleted from the SAN, complete the procedures in this section to expand storage on the Oracle RAC.

Backing Up the Oracle RAC Database

Important: Make sure to save your database backup to a location independent from the Storage Area Network (SAN).

Complete a full backup of the Oracle RAC database. For details about the backup procedures, refer to the ***Explorer Controller Suite 3.0 Backup and Restore User Guide*** (TP-00111).

Add New Disks to the Primary Oracle RAC VM

Complete the following procedures to create new ORADATA and ORABACK disks on the primary RAC VM.

Adding a New ORADATA Disk for the Primary RAC VM

1. Log into the vSphere Web UI. Then select the *primary* RAC VM and click **Edit Settings**.
2. From the bottom of the window, select **New Hard Disk** from the New Device dropdown menu.
3. Click **Add**. The new hard disk is added to the Virtual Hardware list.
4. Click the dropdown arrow to the left of the **New Hard disk** row to view the configuration options for the disk.
5. Modify the following fields.
   a. Enter a value for the disk size. This should be the same disk size as defined for the existing ORADATA disk.
   b. From the **Location** area, click the dropdown arrow and select **Browse**. Then select the **ORADATA** datastore and click **OK**.
   c. From the **Disk Provisioning** area, select **Thick provision eager zeroed**.
   d. From the **Virtual Device Node** section, select the next available virtual device node.

      Important: The Virtual Device Node must be on the Paravirtual SCSI adapter, which is "SCSI(1:x)" not "SCSI (0:x)".

e From the **Disk Mode** area, select I**ndependent - Persistent**.

f Click **OK**. The new disk is reconfigured.

Example Configuration:

Note: The numeric callouts represent the steps in this procedure.

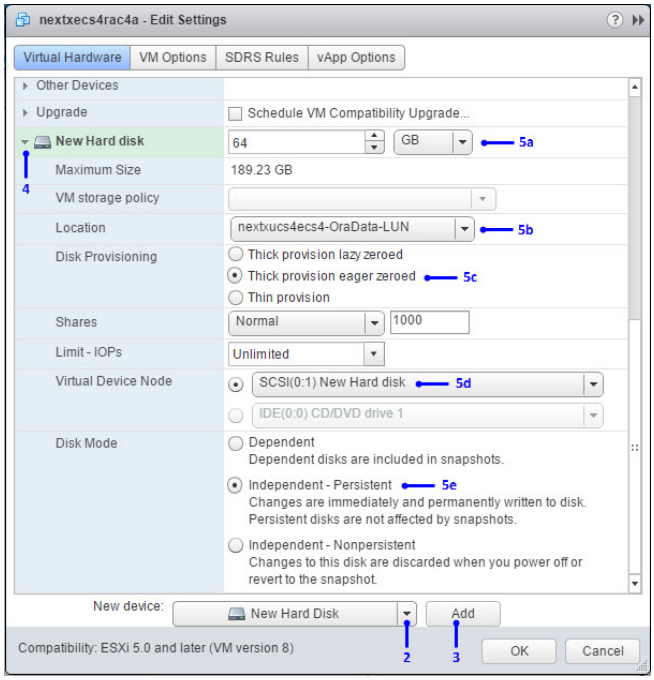

6 Monitor the **Recent Tasks** area to verify that the new disk was reconfigured successfully.

Adding a New ORABACK Disk for the Primary RAC VM

1 Select the *primary* RAC VM again and click Edit Settings. The Edit Settings window displays.

2 From the bottom of the window, select **New Hard Disk** from the New Device dropdown menu.

3 Click **Add**. The new hard disk is added to the Virtual Hardware list.

4 Click the dropdown arrow to the left of the **New Hard disk** row to view the configuration options for the disk.

5 Modify the following fields.

a Enter a value for the disk size. This should be the same disk size as defined for the existing ORABACK disk.

b From the **Location** area, click the dropdown arrow and select **Browse**. Then select the **ORABACK** datastore and click **OK**.

c From the **Disk Provisioning** area, select **Thick provision eager zeroed**.

d From the **Virtual Device Node** section, select the next available virtual device node.

Important: The Virtual Device Node must be on the Paravirtual SCSI adapter, which is "SCSI(1:x)" not "SCSI (0:x)".

e From the **Disk Mode** area, select I**ndependent - Persistent**.

f Click **OK**. The new disk is reconfigured.

Example Configuration:

Note: The numeric callouts represent the steps in this procedure.

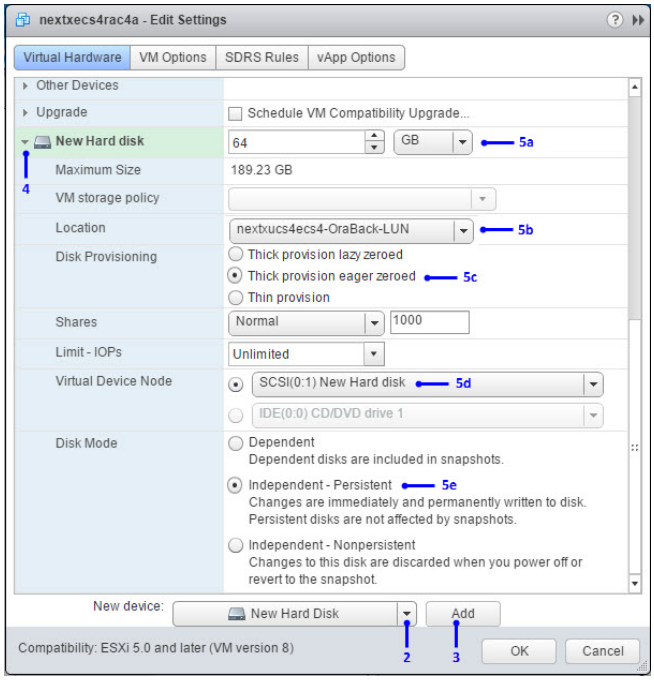

6 Monitor the **Recent Tasks** area to verify that the new disk was reconfigured successfully.

Configuring the New Disks for Clustering/Sharing on the Primary Oracle RAC VM

1 Right-click the primary RAC VM and select **All vCenter Actions > Power > Power Off**.

2 Monitor the **Recent Tasks** area to verify that the task to power off the RAC VM completes successfully.

3 Right-click the primary RAC VM and select **Edit Settings**.

4 Click the **VM Options** tab and then click the dropdown arrow to the left of the **Advanced** row.

5 From the **Configuration Parameters** row, click **Edit Configuration**. The Configuration Parameters window displays.

6 Click **Add Row**. A new row is inserted at the bottom of the table.

7 Enter the following content into each respective text box in the row:

```
scsi1:3.sharing      multi-writer
```

8 Click **Add Row** again and enter the following content into the row.

```
scsi1:4.sharing     multi-writer
```

9 Click **OK** and then click **OK** again.

10 Monitor the **Recent Tasks** area to monitor the reconfiguration of the RAC VM.

11 Right-click the primary RAC VM and select **Power On**.

Add New Disks to the Secondary Oracle RAC VM

Complete the following procedures to create new ORADATA and ORABACK disks on the primary RAC VM.

1 Log into the primary RAC database and change to **oracle** user.

```
[root@nextxecs4rac4a ~]# su oracle
```

2 Enter the following command to source the environment.

```
[oracle@nextxecs4rac4a root]$ . /opt/oracle/CABHE.env
```

3 Enter the following command to verify that the primary RAC VM is up and running.

Note: The output should indicate ONLINE.

```
[oracle@nextxecs4rac4a root]$ srvctl status server -servers
nextxecs4rac4a
Server name: nextxecs4rac4a
Server state: ONLINE
```

4 From the vSphere Web UI, right-click the Secondary RAC VM and select **All vCenter Actions > Power > Power Off**.

5 Monitor the **Recent Tasks** area until the secondary RAC VM is successfully powered off.

Creating a New ORADATA Disk for the Primary RAC VM

1 Log into the vSphere Web UI and select the secondary RAC VM.

2 Right-click the VM and select **Edit Settings**. The Edit Setting window displays.

3 From the bottom of the window, select **Existing Hard Disk** from the New Device dropdown menu. Then click **Add**. The Select File window displays.

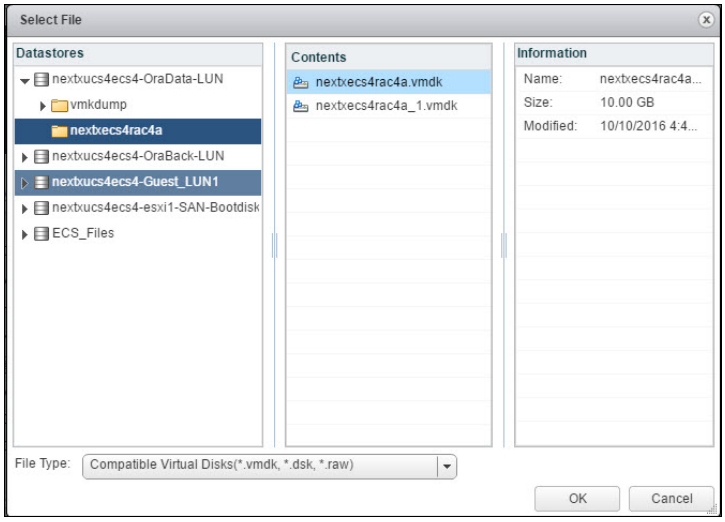

4 Click the **ORADATA** datastore. The Contents area populates.

5 From the Contents area, click the appropriate folder and then select the **vmdk** file for the primary RAC VM. The Information area populates.

6 Click **OK**. The New Hard Hard disk is added to the Edit Settings window.

7 Click **OK** to save the configuration.

8 Monitor the **Recent Tasks** area to verify that the new disk is configured successfully.

Creating a New ORABACK Disk for the Primary RAC VM

1 Log into the vSphere Web UI and select the secondary RAC VM.

2 Right-click the VM and select **Edit Settings**. The Edit Setting window displays.

3 From the bottom of the window, select **Existing Hard Disk** from the New Device dropdown menu. Then click **Add**. The Select File window displays.

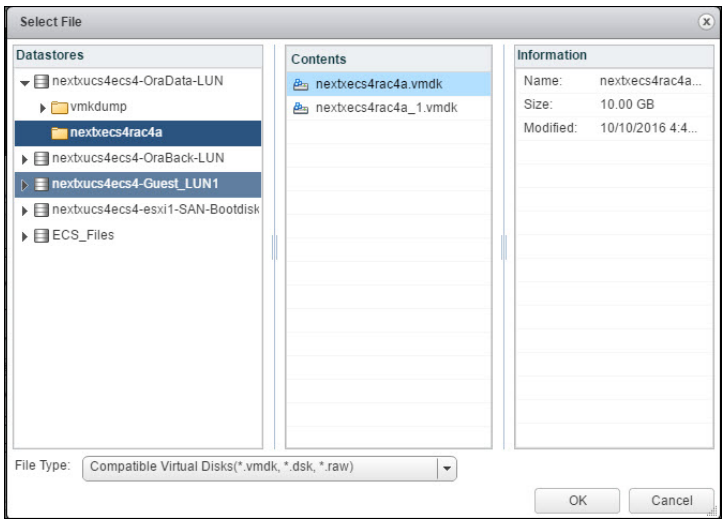

4 Click the **ORABACK** datastore. The Contents area populates.

5 From the Contents area, click the appropriate folder and then select the **vmdk** file for the primary RAC VM. The Information area populates.

6 Click **OK**. The New Hard Hard disk is added to the Edit Settings window.

7 Click **OK** to save the configuration.

8 Monitor the **Recent Tasks** area to verify that the new disk is configured successfully.

Configuring the New Disks for Clustering/Sharing on the Primary Oracle RAC VM

1 Right-click the secondary RAC VM and select **All vCenter Actions > Power > Power Off**.

2 Monitor the **Recent Tasks** area to verify that the task to power off the RAC VM completes successfully.

3 Right-click the secondary RAC VM and select **Edit Settings**. The Edit Setting window displays.

4 Click the **VM Options** tab and then click the dropdown arrow to the left of the **Advanced** row.

5 From the Configuration Parameters row, click **Edit Configuration**. The Configuration Parameters window displays.

6 Click **Add Row**. A new row is inserted at the bottom of the table.

7 Enter the following content into each respective text box in the row:

```
scsi1:3.sharing      multi-writer
```

8 Click **Add Row** again and enter the following content into the row.

```
scsi1:4.sharing      multi-writer
```

9 Click **OK** and then click **OK** again.

10 Monitor the **Recent Tasks** area to monitor the reconfiguration of the RAC VM.

11 Right-click the secondary RAC VM and select **Power On**.

Add the New Disks to Oracle Automatic Storage Management (ASM)

Partitioning the New Disks on the Primary Oracle RAC VM

Complete the following steps to partition the new disks on the primary RAC VM.

1 Login to the *primary* Oracle RAC VM as **root** user. Create a single partition for each new disks that spans the entire disk.

Note: The disks will most likely be sde and sdf.

2 Enter the following command to partition the **/dev/sde** disk.

```
[root@nextxecs4rac4a ~]# fdisk /dev/sde
```

```
Device contains neither a valid DOS partition table, nor Sun, SGI or OSF disklabel
Building a new DOS disklabel with disk identifier 0xdaf546da.
Changes will remain in memory only, until you decide to write them.
After that, of course, the previous content won't be recoverable.

Warning: invalid flag 0x0000 of partition table 4 will be corrected by w(rite)

WARNING: DOS-compatible mode is deprecated. It's strongly recommended to
         switch off the mode (command 'c') and change display units to
         sectors (command 'u').

Command (m for help): n
Command action
   e   extended
   p   primary partition (1-4)
p
Partition number (1-4): 1
First cylinder (1-51200, default 1):
Using default value 1
Last cylinder, +cylinders or +size{K,M,G} (1-51200, default 51200):
Using default value 51200

Command (m for help): w
The partition table has been altered!

Calling ioctl() to re-read partition table.
Syncing disks.
[root@RAC1a rules.d]# fdisk -l /dev/sde

Disk /dev/sde: 53.7 GB, 53687091200 bytes
64 heads, 32 sectors/track, 51200 cylinders
Units = cylinders of 2048 * 512 = 1048576 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk identifier: 0xdaf546da

   Device Boot      Start         End      Blocks   Id  System
/dev/sde1               1       51200    52428784   83  Linux
```

3 Enter the following command to partition the **/dev/sdf** disk.

```
[root@nextxecs4rac4a ~]# fdisk /dev/sdf
```

```
Device contains neither a valid DOS partition table, nor Sun, SGI or OSF disklabel
Building a new DOS disklabel with disk identifier 0x6cb6991f.
Changes will remain in memory only, until you decide to write them.
After that, of course, the previous content won't be recoverable.

Warning: invalid flag 0x0000 of partition table 4 will be corrected by w(rite)

WARNING: DOS-compatible mode is deprecated. It's strongly recommended to
         switch off the mode (command 'c') and change display units to
         sectors (command 'u').

Command (m for help): n
Command action
   e   extended
   p   primary partition (1-4)
p
Partition number (1-4): 1
First cylinder (1-19581, default 1):
Using default value 1
Last cylinder, +cylinders or +size{K,M,G} (1-19581, default 19581):
Using default value 19581

Command (m for help): w
The partition table has been altered!

Calling ioctl() to re-read partition table.
Syncing disks.
[root@RAC1a ~]# fdisk -l /dev/sdf

Disk /dev/sdf: 161.1 GB, 161061273600 bytes
255 heads, 63 sectors/track, 19581 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk identifier: 0x6cb6991f

   Device Boot      Start         End      Blocks   Id  System
/dev/sdf1               1       19581   157284351   83  Linux
```

Adding the Disks to the udev Rules File

Complete the following steps to add the disks to the udev rules file.

1 From the terminal window for the primary RAC VM, enter the following command to determine the SCSI ID for the sde disk.

Note: The highlighted output is the SCSI ID.

```
[root@nextxecs4rac4a ~]# ls -l /dev/disk/by-id/ | grep sde
```

```
lrwxrwxrwx. 1 root root  9 Feb 16 22:53 scsi-36000c291d614342af78ca066b2cd81ec -> ../../sde
lrwxrwxrwx. 1 root root 10 Feb 16 22:53 scsi-36000c291d614342af78ca066b2cd81ec-part1 -> ../../sde1
rwxrwxrwx. 1 root root  9 Feb 16 22:53 wwn-0x6000c291d614342af78ca066b2cd81ec -> ../../sde
lrwxrwxrwx. 1 root root 10 Feb 16 22:53 wwn-0x6000c291d614342af78ca066b2cd81ec-part1 -> ../../sde1
```

2 Enter the following command to determine the SCSI ID for the sdf disk.

Note: The highlighted output is the SCSI ID.

```
[root@nextxecs4rac4a ~]# ls -l /dev/disk/by-id/ | grep sdf
```

```
lrwxrwxrwx. 1 root root  9 Feb 16 22:53 scsi-36000c298134ff70414987890a9047408 -> ../../sdf
lrwxrwxrwx. 1 root root 10 Feb 16 22:53 scsi-36000c298134ff70414987890a9047408-part1 -> ../../sdf1
lrwxrwxrwx. 1 root root  9 Feb 16 22:53 wwn-0x6000c298134ff70414987890a9047408 -> ../../sdf
lrwxrwxrwx. 1 root root 10 Feb 16 22:53 wwn-0x6000c298134ff70414987890a9047408-part1 -> ../../sdf1
```

3 Enter the following command to open the /etc/udev/rules.d/99-oracle-asmdevices.rules file in a text editor.

```
[root@nextxecs4rac4a ~]# vi
/etc/udev/rules.d/99-oracle-asmdevices.rules
```

4 Go to the end of the file and append the following two entries.

Note: Enter each entry on one continuous line.

```
KERNEL=="sd*", SUBSYSTEM=="block", ENV{DEVTYPE}=="disk",
ENV{ID_SERIAL}=="36000c291d614342af78ca066b2cd81ec",
NAME+="asm/ORADATA2", ACTION=="add|change", OWNER="oracle",
GROUP="dba", MODE="0660"
```

```
KERNEL=="sd*", SUBSYSTEM=="block", ENV{DEVTYPE}=="disk",
ENV{ID_SERIAL}=="36000c298134ff70414987890a9047408",
NAME+="asm/ORABACK2", ACTION=="add|change", OWNER="oracle",
GROUP="dba", MODE="0660"
```

5 Using secure copy (scp), copy the **99-oracle-asmdevices.rules** file to the secondary RAC VM.

```
[root@nextxecs4rac4a ~]# scp /etc/udev/rules.d/
99-oracle-asmdevices.rules
root@nextxecs4rac4b:/etc/udev/rules.d/
99-oracle-asmdevices.rules 99-oracle-asmdevices.rules
```

6 Enter the following commands to refresh the udev rules and restart the udev service to put these disks under ASM control.

7 `[root@nextxecs4rac4a ~]# udevadm control --reload-rules`

8 `[root@nextxecs4rac4a ~]#  udevadm trigger --type=devices`

9 `[root@nextxecs4rac4a ~]#  /sbin/start_udev`

Output:

```
Starting udev:                                [  OK  ]
```

10 Repeat Step 9 on the *secondary* RAC VM.

Adding the New Disks to the ASM Diskgroups

Complete the following procedure to add the new disk to the ASM diskgroups.

1 On the *primary* RAC VM, change to **oracle** user.

```
[root@nextxecs4rac4a ~]# su oracle
```

2 Enter the following command to source the ASM environment.

```
[oracle@nextxecs4rac4a root]$ . /opt/oracle/+ASM1.env
```

3 Enter the following command to connect to the database as **sysasm**. The SQL prompt displays.

```
[oracle@nextxecs4rac4a root]$ sqlplus / as sysasm
```

Example:

```
[oracle@nextxecs4rac4a root]$ sqlplus / as sysasm

SQL*Plus: Release 12.1.0.2.0 Production on Thu Jun 1 19:47:46 2017

Copyright (c) 1982, 2014, Oracle.  All rights reserved.


Connected to:
Oracle Database 12c Enterprise Edition Release 12.1.0.2.0 - 64bit Production
With the Real Application Clusters and Automatic Storage Management options
```

4 Enter the following commands to configure sqlplus in a manner that makes the output more readable.

```
SQL> set linesize 200;
SQL> col NAME format a15;
SQL> col LABEL format a25;
SQL> col PATH format a40;
```

5 Enter the following command to display the current state of the ASM diskgroups.

```
SQL> select g.name, d.path, d.os_mb, d.total_mb from
v$asm_diskgroup g, v$asm_disk d where g.group_number =
d.group_number;
```

```
NAME            PATH                                      OS_MB   TOTAL_MB
--------------- ---------------------------------------- ---------- ----------
CRS             /dev/asm/CRS1                              10240      10240
ORADATA         /dev/asm/ORADATA1                         256000     256000
ORABACK         /dev/asm/ORABACK1                        1048576    1048576
```

```
SQL> select name, TOTAL_MB from v$asm_diskgroup;
```

```
NAME             TOTAL_MB
---------------- ----------
ORADATA             256000
ORABACK            1048576
CRS                  10240
```

```
SQL> select group_number, name, TOTAL_MB, FREE_MB from
V$asm_disk_stat;
```

```
GROUP_NUMBER NAME             TOTAL_MB    FREE_MB
------------ --------------- ---------- ----------
           1 CRS_0000             10240       9943
           2 ORADATA_0000        256000     240400
           3 ORABACK_0000       1048576        138
```

6 Enter the following command to show all of the available disks.

```
SQL> select name, PATH, OS_MB, total_mb from v$asm_disk;
```

```
NAME            PATH                           OS_MB  TOTAL_MB
--------------- ------------------------------ ---------- ---------
                /dev/asm/ORADATA2               51200         0
                /dev/asm/oraback_vol-477       153344         0
                /dev/asm/ORABACK2              153600         0
ORABACK_0000    /dev/asm/ORABACK1              153600    153600
ORADATA_0000    /dev/asm/ORADATA1               51200     51200
CRS_0000        /dev/asm/CRS1                   10240     10240
```

7 Enter the following command to add the new disk to the ORADATA diskgroup.

```
SQL> alter diskgroup ORADATA add disk '/dev/asm/ORADATA2';
```

8 Enter the following command to add the new disk to the ORABACK diskgroup.

```
SQL> alter diskgroup ORABACK add disk '/dev/asm/ORABACK2';
```

Note: The resizing of the disks will occur in the background. It is safe to continue with the next command to allow both disks to resize simultaneously.

9 Enter the following command to check the status of resizing at any point.

```
SQL> select * from v$asm_operation;
```

10 Enter the following command to verify that the diskgroups were expanded successfully.

```
SQL> select g.name, d.path, d.os_mb, d.total_mb from
v$asm_diskgroup g, v$asm_disk d where g.group_number =
d.group_number;
```

Example Output:

```
NAME            PATH                            OS_MB  TOTAL_MB
--------------- ------------------------------ ---------- ----------
CRS             /dev/asm/CRS1                   10240     10240
ORADATA         /dev/asm/ORADATA2               51200     51200
ORADATA         /dev/asm/ORADATA1               51200     51200
ORABACK         /dev/asm/ORABACK2              153600    153600
ORABACK         /dev/asm/ORABACK1              153600    153600


SQL> select name, TOTAL_MB from v$asm_diskgroup;

NAME             TOTAL_MB
--------------- ----------
ORADATA             102400
ORABACK             307200
CRS                  10240

SQL> select group_number, name, TOTAL_MB, FREE_MB from V$asm_disk_stat;

GROUP_NUMBER NAME           TOTAL_MB    FREE_MB
------------ --------------- ---------- ----------
           1 CRS_0000          10240       9943
           2 ORABACK_0000     153600       3464
           3 ORADATA_0000      51200      13563
           2 ORABACK_0001     153600     150286
           3 ORADATA_0001      51200      42802
```

Note: In a short time the existing data will be spread evenly across the disks in each disk group.

Example: ORABACK disk group still rebalancing

```
GROUP_NUMBER NAME           TOTAL_MB    FREE_MB
------------ --------------- ---------- ----------
           1 CRS_0000          10240       9943
           2 ORABACK_0000     153600      22688
           3 ORADATA_0000      51200      28172
           2 ORABACK_0001     153600     131062
           3 ORADATA_0001      51200      28193
```

6

Verify ECS Functionality

This chapter includes the steps to verify the functionality of the ECS 3.0 system post installation or post migration.

In This Chapter

Verifying ECS Functionality

Once you have installed or migrated to ECS 3.0, you should verify the functionality of the ECS system. The following features should be tested from the ECS Web UI.

Note: When executing these tests, go to the **/opt/jboss-as/standalone/log** directory to monitor the appropriate log.

- **Operations and Administration (OAM)**
  - **Command Execution**: execute system commands on the remote EC/DTACS system
  - **Backup and Restore**: verify a software backup
- **Regional Provisioning (RPS)**
  - **Export from EC Device**: test that you can export data from the EC database and generate the export data file
  - **Import XML Template**: generate an apply (ingest) file from an XLS template based on the selected provisioning element (PE), and load the generated ingest file
  - **Apply to EC Device(s)**: apply (ingest) a request from the RPS to the RPO (Regional Provisioning Orchestrator), a service running on the EC. The RPO should process the apply request and then fetch the apply file using SFTP. It should then parse the file and provision the data onto the EC
- **CPE Management**
  - **Manage CPE**: verify that you can query a CPE
  - **Batch Management**: test that you can perform a batch install

7

Customer Information

If You Have Questions

If you have technical questions, call Cisco Services for assistance. Follow the menu options to speak with a service engineer.

Access your company's extranet site to view or order additional technical publications. For accessing instructions, contact the representative who handles your account. Check your extranet site often as the information is updated frequently.

A

RAC Install Using ESXi 6.0u1 or Later

This appendix includes the procedures specific to ESXi 6.0u1 or later to install the Oracle RAC in your NextX system.

Note: If you are using ESXi 6.0 or earlier, do not execute any procedures in this appendix.

In This Appendix

Modifying the RAC Configuration File

Complete the following procedure to create and update the RAC configuration file. This file is used during installation of the RAC.

1 From your local PC, go to the directory where you downloaded the **oracle-rac-ee-3.0.3_20161129.tar** file.

Note: This should already be downloaded to your local PC. If it has not yet been downloaded, refer to ***Software Requirements*** (on page 3).

2 Untar the **oracle-rac-ee-3.0.3_20161129.tar** file. The following files are extracted:

- Prepare-OracleRAC_VM_RHEL6_v_2_07A.ps1
- RAC_12c_rpms_RHEL6.x_EE.iso
- RAC_8CPU_64GBRAM_80GBHDD_RHEL6U8/
- RAC_8CPU_64GBRAM_80GBHDD_RHEL6U8/RAC_8CPU_64GBRAM_80GBHDD_RHEL6U8.mf
- RAC_8CPU_64GBRAM_80GBHDD_RHEL6U8/RAC_8CPU_64GBRAM_80GBHDD_RHEL6U8-disk1.vmdk
- RAC_8CPU_64GBRAM_80GBHDD_RHEL6U8/RAC_8CPU_64GBRAM_80GBHDD_RHEL6U8.ovf
- RAC_addon_scripts_RHEL6.x_EE_v2.0.8.iso
- rac.cfg.geo-active.sample
- rac.cfg.geo-standby.sample
- rac.cfg.non-geo.sample

3 Copy the **rac.cfg.non-geo.sample** file to a new file name that reflects the Oracle RAC system being built.

Example: Copy rac.cfg.non-geo.sample nextxecs3rac.cfg

4 Open the new file (i.e. nextxecs3rac.cfg) file in a Wordpad (or vi editor on a Mac or Linux system) and update the following fields:

Important: Enter a value for each field listed below even if the field states "leave empty".

- **HOST_GROUP_NAME** — The base of the host names of the Oracle RAC VMs

  Example: nextxecs3

- **PRIMARY_HOST** — Host name of primary RAC VM

  Example: nextxecs3rac1

- **SECONDARY_HOST** — Host name of secondary RAC VM

  Example: nextxecs3rac2

- **PRIMARY_IP —** IP address (eth0) to access the primary RAC VM
- **SECONDARY_IP —** IP address (eth0) to access the secondary RAC VM
- **SCAN_ALIAS —** Host name for the SCAN interfaces

  Example: nextxecs3scan
- **SCAN_IP_1 —** Primary IP address used by services to access the database
- **SCAN_IP_2 —**Secondary IP address used by services to access the database
- **VIP_ALIAS_1 —** Host name alias for VIP interface 1
- **VIP_ALIAS_2 —** Host name alias for VIP interface 2
- **VIP_IP_1 —** Virtual IP address for VIP interface 1

  Example: nextxecs3vip1
- **VIP_IP_2 —** Virtual IP address for VIP interface 2

  Example: nextxecs3vip2
- **MGMT_ALIAS_1 —** Host name alias for primary private/cluster interface

  Example: nextxecs3mgmt1
- **MGMT_ALIAS_2 —** Host name alias for secondary private/cluster interface

  Example: nextxecs3mgmt2
- **MGMT_IP_1 —** High Availability private/cluster IP address used to communicate between the RAC hosts
- **MGMT_IP_2 —** High Availability private/cluster IP address used to communicate between the RAC hosts
- **ETH0_GW —** Default Gateway IP address
- **ETH0_MASK —** Netmask of eth0 interface
- **ETH1_MASK —** Netmask of eth1 interface
- **DATA_DISK_SIZE —** 512 (default value). The disk size should be between 15 GB and 500 GB depending on the number of DHCTs that will be supported.
- **BACK_DISK_SIZE —** 1024 (default value)
- **ORACLESID —** CABHE
- **BUILD_INSTANCE —** Maintain default value, YES
- **STARTING_INDEX —** Maintain default value, 1.
- **ORACLE_CLUSTER_NAME —** Maintain default value, ora-cluster
- **AUTO_SIZE_ORACLE_MEM —** Automatically sizes the memory; maintain default value, YES
- **ENABLE_RMAN —** Maintain default value, YES
- **INSTALL_DG_RPM —** Maintain default value, NO

- **NTP_PRIMARY —** IP address of primary NTP server
- **NTP_BACKUP —** IP address of secondary NTP server
- **NTP_ENABLE —** Default is YES; change to NO if no NTP services will be enabled.

  Note: Cisco strongly recommends using NTP.

5 Save and close the file.

6 Copy the file (i.e. nextxecs3rac.cfg) file to **rac.cfg**.

Note: When entering the file name (rac.cfg) in Wordpad, you MUST enclose it in quotes ("). This prevents WordPad from attaching a suffix to the file name (for example, rac.cfg.doc, or rac.cfg.txt).

Deploying the Oracle RAC VMs

Complete the following procedure to deploy two Oracle RAC VMs for your NextX ECS 3.0 system.

1 Via a Web browser, log into your vCenter 6.0u1 or later vSphere Web client.

2 Click **Hosts and Clusters** to view the datacenters and ESXi hosts.

3 Right-click the appropriate ESXi host where you want to deploy the *primary* RAC VM and select **Deploy OVF Template**. The Deploy OVF Template window displays.

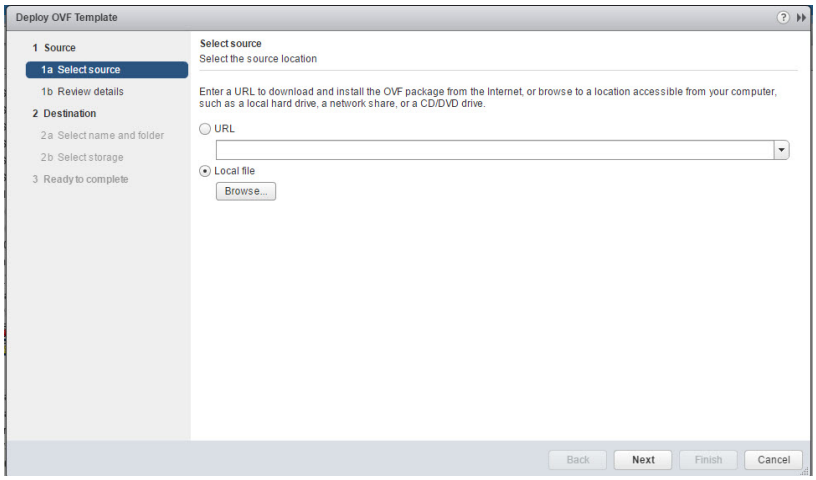

4 Select Local file and click **Next**. The Open window displays.

5 Navigate to the location where you saved the **RAC_8CPU_64GBRAM_80GBHDD_RHEL6U8/RAC_8CPU_64GBRAM_80GBHDD_RHEL6U8.ovf** file.

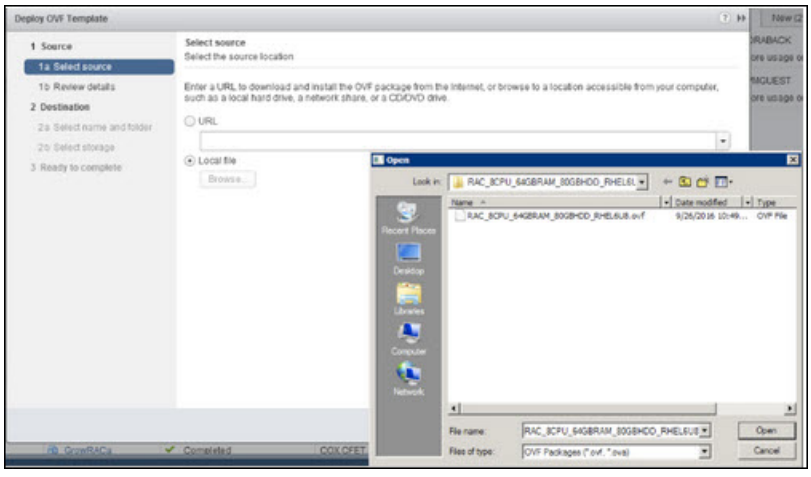

6 Select the OVF file and click **Open**. Then click **Next**. The Select name and folder window opens.

7 From the **Name** text box, enter a name to describe the VM (in this example, we will use RAC1 for the primary RAC VM and RAC2 for the secondary RAC VM).

8 From the **Select a folder or datacenter** area, select the appropriate datacenter.

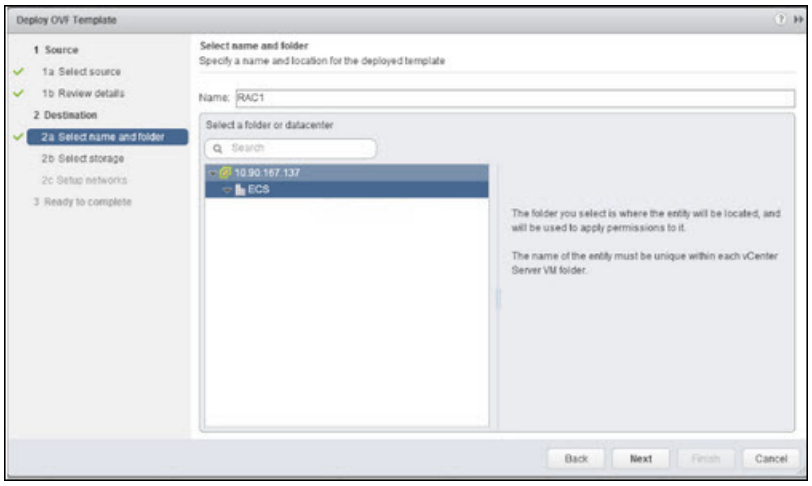

9 Click **Next**. The Select storage window displays.

10 From the **Select virtual disk format** dropdown menu, select **Thick Provision Lazy Zeroed**.

11 From the **VM Storage Policy** dropdown menu, select **Datastore Default**.

12 From the **available datastores** area, select the **VMGUEST** datastore.

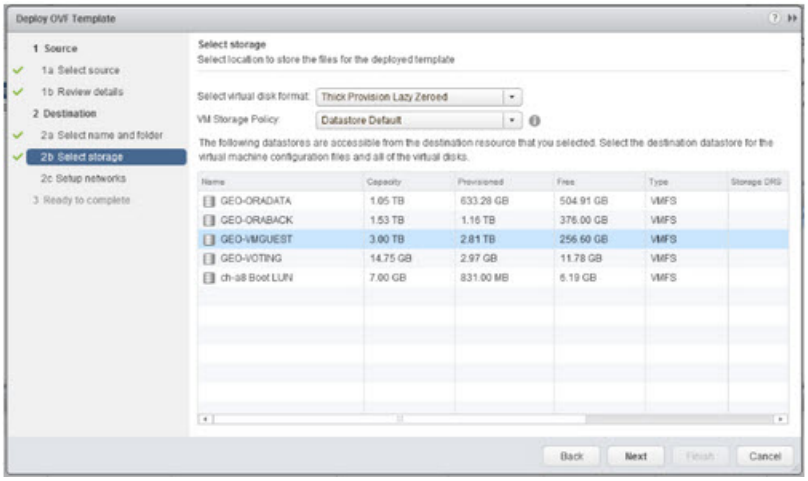

13 Click **Next**. The Setup networks window displays.

14 Select the appropriate network.

Notes:

- Select the **eth0** interface as this is used to host the database.
- The service IP addresses, virtual IP addresses (VIPs) and SCAN IP addresses will use this interface.
- Additional network interfaces are created later.

15 Click **Next**. The Ready to Complete window displays.

16 Click **Finish**.

Important: Do not power on the VM yet.

17 Repeat Steps 3 through 16 to deploy the *secondary* RAC VM.

Important: Remember to select a different ESXi host in Step 3 to deploy the secondary RAC VM.

Reconfiguring the Oracle RAC VMs

Adding Network Adapters to the RAC VMs

Complete the following procedure to add additional network adapters to both RAC VMs.

1 From the vSphere Web client, right-click the *primary* RAC VM and select **Edit Settings**. The Edit Settings window displays.

2 From the **New device** dropdown menu (bottom of the window), select **Network** and then click **Add**. A new Network adapter is added to the list of Virtual Hardware.

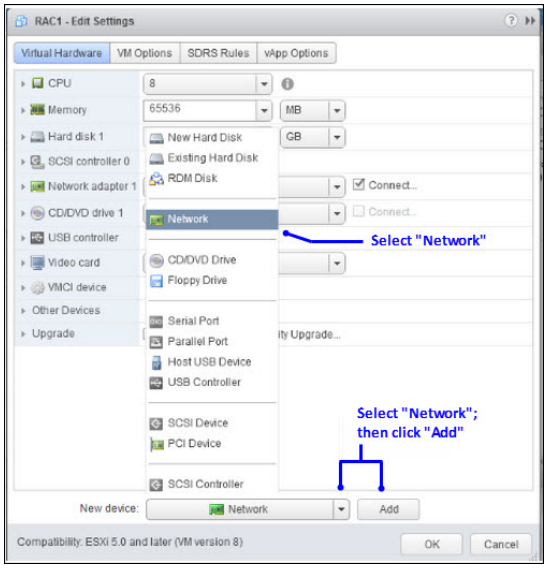

3 Click the dropdown menu to select your Private Oracle VLAN (virtual local area network).

Note: This interface is the eth1 interface.

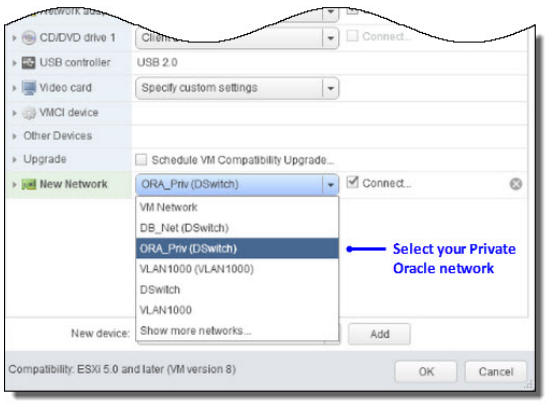

4 Is Network adapter 1 (eth0) on a restricted network?

- If **yes**, go to the next step.
- If **no**, go to Step 6.

5 Repeat Steps 2 through 3 to add a third interface for management access to the RAC VM.

Note: Remember to select the network for management access when repeating Step 3.

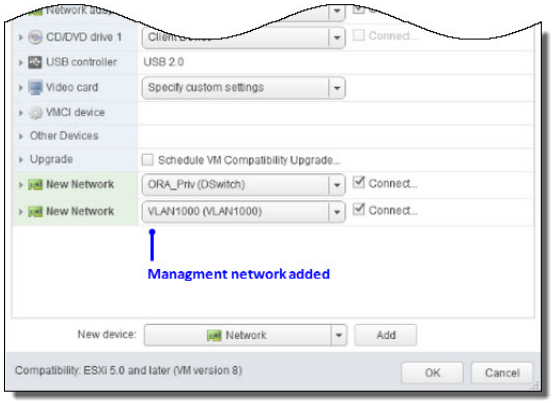

6 Click **OK** to reconfigure the VM.

7 Monitor the **Recent Tasks** area to verify that the VM was successfully reconfigured.

8 Repeat Steps 1 through 7 to configure the networks for the *secondary* RAC VM.

Adding Shared Disks

Complete the following procedure only on the *primary* RAC VM to add shared disks to each RAC VM.

1. Select the *primary* RAC VM and click **Edit Settings**. The Edit Settings window appears.
2. Click the **VM Options** tab.
3. From the left column, click the down arrow, ▾, next to **Advanced**. The advanced settings are displayed.

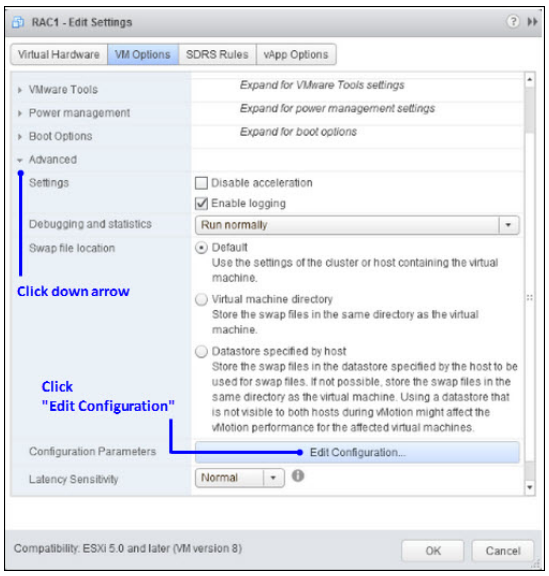

4. Click **Edit Configuration**. The Configuration Parameters window displays.
5. Click **Add Row**. A new row is added to the table.
6. Click in the **Name** column for the new row and type **disk.EnableUUID**.
7. Press the **Tab** key and in the **Value** column, type **TRUE**.

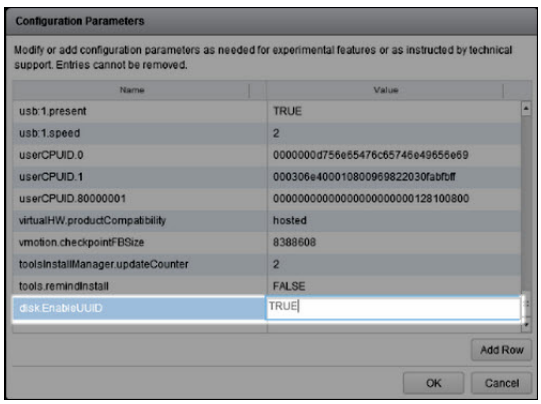

8. Click **OK**. You are returned to the **Edit Settings > VM Options** window.
9. Click **OK** again.

10 Monitor the **Recent Tasks** area to verify that the VM is successfully reconfigured.

Adding a New SCSI Controller to Each RAC VM

Complete the following procedure to add a second SCSI controller to each RAC VM.

Note: All of the shared disks will be attached to this controller.

1 Right-click the *primary* RAC VM and select **Edit Settings**.

2 Click the **New device** dropdown menu and select **SCSI Controller**. The Add button becomes active.

3 Click **Add**. A new SCSI controller is added to the list of Virtual Devices.

4 From the left column in the window, click **New SCSI controller**. Settings for the controller appear below the entry.

5 If needed, update the fields to the following values.

- **SCSI Bus Sharing** None
- **Change Type** VMware Paravirtual

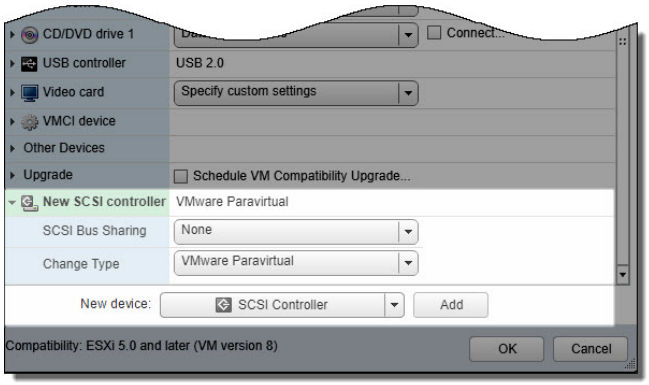

6 Click **OK**. The Edit Settings window closes.

7 Monitor the **Recent Tasks** area until the Status indicates **Completed**.

8 Repeat Steps 1 through 7 on the *secondary* RAC VM.

Creating the Shared Disks on the Primary RAC VM

Complete the following steps to create the shared disks on the *primary* RAC VM.

Note: These disks will be added to the *secondary* RAC VM in the next procedure.

1 From the vSphere Web UI, right-click the *primary* RAC VM and select **Edit Settings**.

2 From the **New device** dropdown menu, select **New Hard Disk** and then click **Add**. The new hard disk is added to the list of Virtual Hardware.

Note: This disk will be the Voting Disk.

3 Modify the disk size to **10 GB**.

Note: This disk size is always 10 GB regardless of the size of the other database disks.

4 Click on the **New Hard Disk** entry to display additional settings.

5 Update the following fields as described below.

Field	Value
Disk Provisioning	Thick provision eager zeroed
Location	Click the dropdown arrow and select the Voting or Oradata datastore. **Note**: If the datastore is not listed, click **Browse** to navigate to the appropriate datastore.
Sharing	Multi-writer
Virtual Device Node	SCSI controller 1 **Note**: The adjacent dropdown should automatically default to SCSI(1:0).
Disk Mode	Independent - Persistent

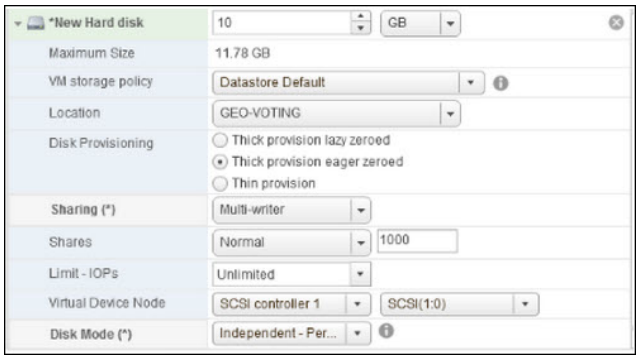

6 From the **New device** dropdown where **New Hard Disk** is selected, click **Add** again. A second hard disk is added to the list of Virtual Devices.

Note: This disk is the Oradata disk.

7 Modify the disk size to the Oradata-recommended size for the installation type.

Note: The disk size should be between 15 GB and 500 GB depending on the number of DHCTs that will be supported.

8 Click on the **New Hard Disk** entry to display additional settings.

9 Update the following fields as described below.

Field	Value
Disk Provisioning	Thick provision eager zeroed
Location	Click the dropdown arrow and select the Oradata datastore. **Note**: If the datastore is not listed, click **Browse** to navigate to the appropriate datastore.
Sharing	Multi-writer
Virtual Device Node	SCSI controller 1 **Note**: The adjacent dropdown should automatically default to SCSI(1:1).
Disk Mode	Independent - Persistent

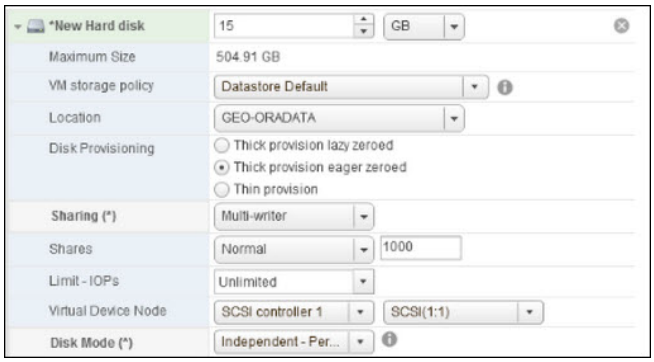

10 From the **New device** dropdown where **New Hard Disk** is selected, click **Add** again. A third hard disk is added to the list of Virtual Devices.

Note: This disk is the Oraback disk.

11 Modify the disk size to the Oraback-recommended size for the installation type.

Note: Typically, this disk size is about three times the size of the Oradata disk.

12 Click on the **New Hard Disk** entry to display additional settings.

13 Update the following fields as described below.

Field	Value
Disk Provisioning	Thick provision eager zeroed
Location	Click the dropdown arrow and select the Oraback datastore. **Note**: If the datastore is not listed, click **Browse** to navigate to the appropriate datastore.
Sharing	Multi-writer

Field	Value
Virtual Device Node	SCSI controller 1 **Note**: The adjacent dropdown should automatically default to SCSI(1:2).
Disk Mode	Independent - Persistent

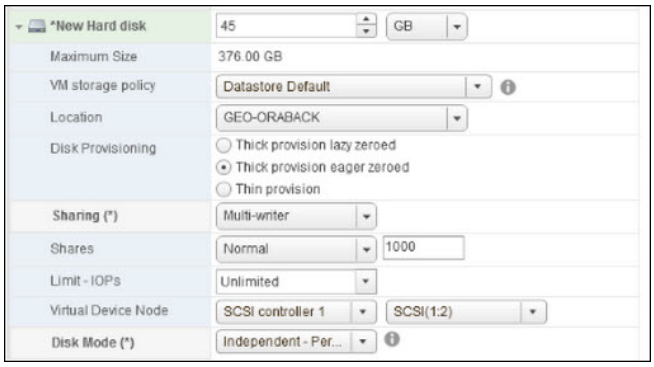

14 Click **OK**. The Edit Settings window closes.

15 Monitor the **Recent Tasks** area until the status for the task indicates **Completed**.

Note: This may take over an hour to complete as it depends on the disks sizes and the shared storage performance.

16 When the status for the task indicates **Completed**, go to the next section.

Adding the Shared Disks to the Secondary RAC VM

Important: Do not start this procedure until the shared disks you added to the primary RAC VM are successfully configured.

Complete the following steps to add the shared disks from the primary RAC VM to the *secondary* RAC VM.

1 From the vSphere Web UI, right-click the *secondary* RAC VM and select **Edit Settings**.

2 From the **New device** dropdown menu, select **Existing Hard Disk** and then click **Add**. The Select File window displays and lists the datastores.

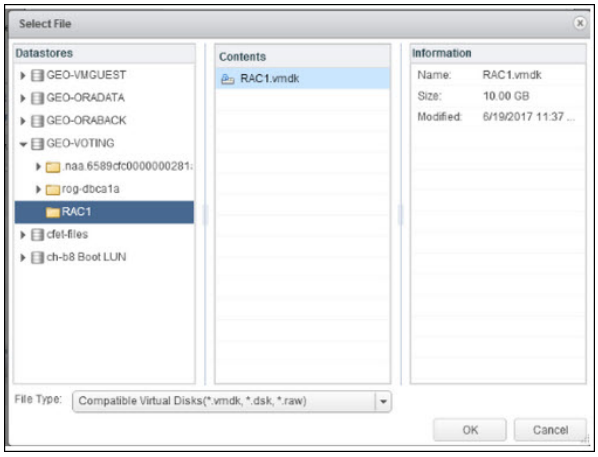

3 Navigate to the appropriate folder where the **Voting** disk for the *primary* RAC VM resides.

4 Select the **Voting** disk and then click **OK**.

5 Configure the disk identical to the settings shown in the image of Step 5 of the previous procedure.

6 Repeat Steps 2 through 4 for the **Oradata** disk and configure the disk identical to the settings shown in the image in Step 9 of the last procedure.

7 Repeat Steps 2 through 4 for the **Oraback** disk and configure the disk identical to the settings shown in the image in Step 13 of the last procedure.

8 Click **OK**. The Edit Settings window closes.

9 Monitor the **Recent Tasks** area until the status for the task indicates **Completed**.

Note: This may take over an hour to complete as it depends on the disks sizes and the shared storage performance.

10 When the status for the task indicates **Completed**, go to the next section.

Configuring the Linux OS on the RAC VMs

Important: The RAC_addon_scripts_RHEL6.x_EE_v2.0.8.iso is in the folder where you extracted the tar file on your local PC.

Complete the following steps to configure the Linux operating system on each RAC VM.

1 From the vSphere Web UI, right-click the *primary* RAC VM and select **Power > Power On**.

2 Select the primary RAC VM again and then click the **Summary** tab.

3 Click **Launch Remote Console**. The VMware Remote Console window displays.

Note: If you have not yet installed the Remote Console, click **Download Remote Console** to install it. Once it is installed, repeat Step 3.

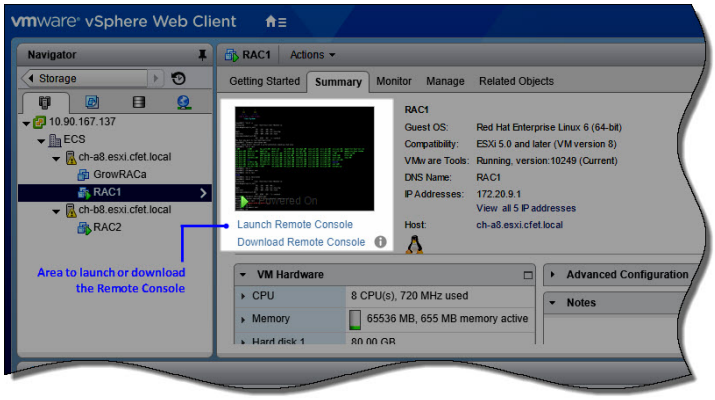

4 Click the **CD/DVD drive 1** icon, , and select **Connect to Disk Image File (iso)**. The Choose Image window opens.

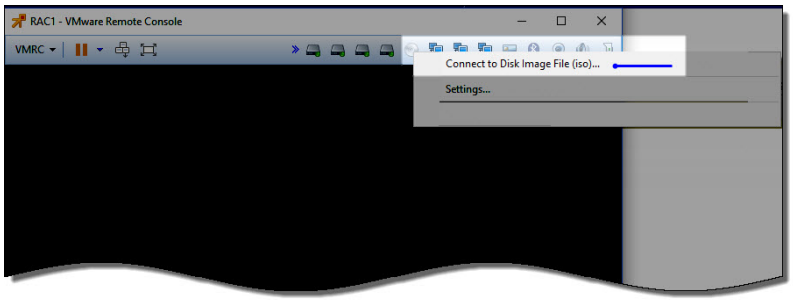

5 Navigate to the **RAC_addon_scripts_RHEL6.x_EE_v2.0.8.iso** file and select **Open**.

6 Log into the RAC VM with the following credentials:

Username:root

Password: vgh12345678

7 Execute the following commands to mount the ISO and copy the scripts to the RAC VM.

```
[root@rac1 ~]# mount /dev/cdrom /mnt
[root@rac1 ~]# cp -f /mnt/*.sh /opt/nds/custom_tools/utils
```

8 Enter the following command to change the access permissions to **0700**.

```
[root@rac1 ~]# chmod 0700 /opt/nds/custom_tools/utils/*.sh
```

9 Enter the following command to unmount the ISO.

```
[root@rac1 ~]# umount /mnt
```

10 From the Remote Console, click the **CD/DVD drive 1** icon and select the **Disconnect** option.

11 Repeat Steps 1 through 10 for the *secondary* RAC VM.

Run the Network Configuration Scripts on Each RAC VM

Complete the following procedure to run the network configuration scripts on both RAC VMs.

1 Enter the following command run the RAC network script on the *primary* RAC VM.

Command Syntax:

```
/opt/nds/custom_tools/utils/rac_network.sh –H [hostname] –I
[eth0 IP] –N [eth0 Netmask] –G [eth0 Gateway] –I [eth1 IP] –n
[eth1 Netmask]
```

Example:

```
/opt/nds/custom_tools/utils/rac_network.sh -H RAC1 -I
172.20.9.1 -N 255.255.255.0 -G 172.20.9.254 -i 192.168.1.1 -n
255.255.255.0
```

2 Repeat Step 1 on the *secondary* RAC VM.

3 Copy the **rac.cfg** file from your local computer to the **/opt/nds/custom_tools/etc/** directory *each* RAC VM.

Note: This rac.cfg file was created and saved on your local PC in ***Modifying the RAC Configuration File*** (on page 144).

Command Syntax:

```
scp rac.cfg root@[RAC1 IP address]:/opt/nds/custom_tools/etc/
scp rac.cfg root@[RAC2 IP address]:/opt/nds/custom_tools/etc/
```

Examples:

```
scp rac.cfg root@172.20.9.1:/opt/nds/custom_tools/etc/
scp rac.cfg root@172.20.9.2:/opt/nds/custom_tools/etc/
```

4 Was the **rac.cfg** file created on a Windows machine?

- If **yes**, go to the next step.
- If **no**, go to the next section, ***Defining the Shared Disks***.

5 Execute the following command on each RAC VM to remove DOS carriage returns from the **rac.cfg** file.

```
[root@rac1 ~]# dos2unix /opt/nds/custom_tools/etc/rac.cfg
```

6 Go to the next section in this appendix.

Defining the Shared Disks

Complete the following procedure to defining the shared disks.

1 Enter the following command on the *primary* RAC VM to configure the shared disks.

```
[root@rac1 ~]# /opt/nds/custom_tools/utils/rac_fdisk.sh
-S -t P
```

2 Enter the following command on the `secondary` RAC VM to configure the shared disks.

```
[root@rac2 ~]# /opt/nds/custom_tools/utils/rac_fdisk.sh
-S -t S
```

Important: Notice that the last option in this command, S, is different than the last option in Step 1.

Configuring the Oracle User

Complete the following steps to create an oracle user on both the primary and secondary RAC VMs.

1. Enter the following command on the *primary* RAC VM to create the **oracle** user account.

   ```
   [root@rac1 ~]# /opt/nds/custom_tools/utils/rac_orauser.sh
   ```

2. Repeat Step 1 on the *secondary* RAC VM.

Configuring Password-less SSH Between the Root and Oracle User Accounts

1. Enter the following command on the *primary* RAC VM to configure password-less SSH access between the root and oracle user accounts.

   ```
   [root@rac1 ~]# /opt/nds/custom_tools/utils/rac_Vsshkeys.sh
   ```

2. Repeat Step 1 on the *secondary* RAC VM.

3. Execute the following commands on the *primary* RAC VM to configure password-less SSH between the two RAC VMs.

   Note: Replace [RAC2_Hostname] with the actual hostname for your secondary RAC VM.

   ```
   [root@rac1 ~]# cat /root/.ssh/authorized_keys >
   /tmp/authorized_keys
   [root@rac1 ~]# ssh root@[RAC2_Hostname] cat
   /root/.ssh/authorized_keys >> /tmp/authorized_keys
   [root@rac1 ~]# cat /home/oracle/.ssh/authorized_keys | sed -e
   "s/root@/oracle@/g" >> /tmp/authorized_keys
   [root@rac1 ~]# ssh root@[RAC2_Hostname] cat
   /home/oracle/.ssh/authorized_keys | sed -e "s/root@/oracle@/g"
   >> tmp/authorized_keys

   [root@rac1 ~]# cp /tmp/authorized_keys /root/.ssh/
   [root@rac1 ~]# cp /tmp/authorized_keys /home/oracle/.ssh/

   [root@rac1 ~]# scp -Crp /tmp/authorized_keys
   root@[RAC2_Hostname]:/root/.ssh/
   [root@rac1 ~]# scp -Crp /tmp/authorized_keys
   root@[RAC2_Hostname]:/home/oracle/.ssh/

   [root@rac1 ~]# cp /root/.ssh/known_hosts /home/oracle/.ssh/
   [root@rac1 ~]# scp /root/.ssh/known_hosts root@[RAC2_Hostname]
   :/root/.ssh/known_hosts
   [root@rac1 ~]# scp /root/.ssh/known_hosts
   root@[RAC2_Hostname]:/home/oracle/.ssh/
   ```

   Note: Do ***not*** repeat Step 3 on the secondary RAC VM. It is only required to be executed on the primary RAC VM.

Configuring NTP on Each RAC VM

1. Execute the following command on the *primary* RAC VM to configure NTP.

   ```
   /opt/nds/custom_tools/utils/rac_ntp.sh -c
   /opt/nds/custom_tools/etc/rac.cfg -S
   ```

2. Repeat Step 1 on the *secondary* RAC VM.

Installing the Oracle RAC Software

Important: Complete this procedure only on the *primary* RAC VM.

Complete the following steps to mount the RAC_12c_rpms_RHEL6.x_EE.iso file from your local computer.

Note: The RAC_12c_rpms_RHEL6.x_EE.iso is in the folder where you extracted the tar file on your local PC.

1. From the **Remote Console**, click the **CD/DVD drive 1** icon, , and select **Connect to Disk Image File (iso)**. The Choose Image window opens.
2. Navigate to the **RAC_12c_rpms_RHEL6.x_EE.iso** file and select **Open**.
3. Execute the following commands to mount the ISO and copy the scripts to the RAC VM.

   ```
   [root@rac1 ~]# mount /dev/cdrom /mnt
   [root@rac1 ~]# cp -f /mnt/*.rpm /opt/nds/custom_tools/utils
   ```

4. Enter the following command to change the access permissions to **0700**.

   ```
   [root@rac1 ~]# chmod 0700 /opt/nds/custom_tools/utils/*.rpm
   ```

5. Enter the following command to install the rpm packages.

   ```
   [root@rac1 ~]# /opt/nds/custom_tools/utils/rac_rpm_install.sh
   CABHE
   ```

6. Click the **CD/DVD drive 1** icon, , and select the **Disconnect** option to unmount the ISO image.
7. Repeat Steps 1 through 6 in the remote window of the *secondary* RAC VM.
8. Execute the following command on the *primary* RAC VM only.

   Note: This script will take around an hour to run. Do ***not*** repeat this step on the *secondary* RAC VM.

   ```
   [root@rac1 ~]# /opt/nds/custom_tools/utils/rac_run_launcher.sh
   CABHE
   ```

9. When the installation completes, go to ***RAC Installation Verification Procedures*** (on page 21).

B

Regionalize ECs and DTACS Servers to the ECS

This appendix provides the procedures to configure and regionalize EC SR 8.0 and DTACS SR 5.0 servers to an ECS system, as well as additional features such as unregionalizing or deleting a registration.

Important: The procedures are nearly identical for registering ECs and DTACS servers. Where steps differ, instructions are provided for both an EC and a DTACS server.

In This Appendix

Configuring the Consul Configuration File

Complete the following procedures to configure the Consul configuration file, config.json, on the EC/DTACS that will be regionalized to the SR 3.0 ECS system.

1 Login to the EC/DTACS as **admin** user.

2 Change to **root** user.

```
[admin@vodwater ~]$ sudo -i
```

3 Enter the following command to verify that the Consul package is installed.

```
[root@vodwater ~]# rpm -qa | grep -i consul
```

4 Is the Consul package present?

- If **yes**, go to Step 5.
- If **no**, enter the following command and then repeat Step 3.

```
[root@vodwater ~]# yum install cisco-vcs-consul
[root@vodwater ~]# rpm -qa | grep -i consul
```

5 Enter the following command to change to the **/etc/consul/** directory.

```
[root@vodwater consul]# cd /etc/consul
```

6 Enter the following command to copy the **client.json.template** to a new file named **config.json**.

```
[root@vodwater consul]# cp -p client.json.template config.json
```

7 Open the **config.json** file in a text editor.

8 Update the following fields with values specific to your EC or DTACS client and specific to the IP Addresses of the three Consul nodes on the ECS system.

Note: Only edit these four lines.

- <client_ip>
- <server_ip1>
- <server_ip2>
- <server_ip3>

Example File:

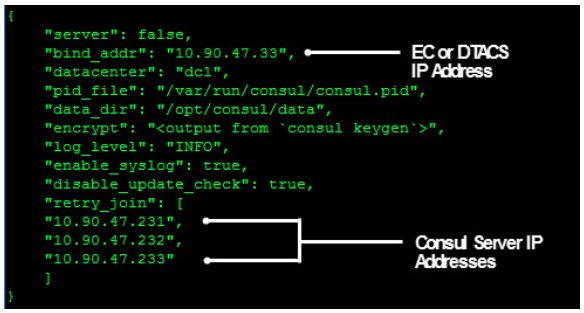

9 Save and close the file.

10 Enter the following command to change the permissions of the config.json file to **0600**.

```
[root@vodwater consul]# chmod 0660 config.json
```

11 Start the consul service.

```
[root@vodwater consul]# service consul start
```

12 Verify that the consul service started.

```
[root@vodwater consul]# service consul status
```

13 Run the following command to stream logs from a Consul agent and verify that it is running successfully.

```
[root@vodwater consul]# consul monitor
```

14 Enter the following commands to verify that the EC is now a member of the Consul cluster.

```
[root@vodwater consul]# consul members
```

```
Node                 Address              Status  Type    Build  Protocol  DC
consul-condor2       10.90.47.26:8301     alive   server  0.7.2  2         dc1
vcs-ecs30-condor2    10.90.47.9:8301      alive   client  0.7.2  2         dc1
vcsconsole-condor2   10.90.47.18:8301     alive   client  0.7.2  2         dc1
vodwater             10.90.45.181:8301    alive   client  0.7.2  2         dc1
```

15 Enter the following command to enable the Consul service to start on boot up.

```
[root@vodwater consul]# chkconfig consul on
```

Enabling HTTPS on an EC/DTACS Server

Important:

- SSL certificates should be installed on the EC or DTACS server, as well as on all ECS nodes.
- If the EC was built via an EC migration, manually compare the .orig files included in the /disk1/keyfiles_staging/etc/apache2/user-conf directory with the Apache .conf files located in the /etc/httpd/user-conf/ directory. Update the .conf files as needed.
- In this procedure, Cisco Labcerts are used.

Complete the following procedures to enable HTTPS on an EC or DTACS server that includes a new SR 8.0 installation.

1 On the EC, change to the **/opt/cisco/vcs** directory.

```
[root@vodwater consul]# cd /opt/cisco/vcs
```

2 Modify the access permissions to **644** and the ownership to **root:root**.

```
[root@vodwater opt]# chmod 644 security.properties
[root@vodwater opt]# chown root:root security.properties
```

3 Open the **security.properties** file in a text editor and update it, as needed, to duplicate the following content.

Note: The fields to typically update only include TrustStoreFile, TrustStorePasswd, KeyStoreFile and KeyStorePasswd.

Example Output:

```
SSLType=TwoWayAuthentication
TrustStoreFile=/etc/pki/tls/vodwater.domainTruststore.jks
TrustStorePasswd=2g3n3r!c
KeyStoreFile=/etc/pki/tls/vodwater.domainKeystore.jks
KeyStorePasswd=2g3n3r!c
RestClientSSLProtocols=TLSv1.1,TLSv1.2
ProtocolVersion=TLSv1.1
ignoreHttpsHost=true
VerifyHostname=false
BasicAuthentication=true
BasicUsername=restful
BasicPasswd=conductor
ignoreHttpsHost=true
```

4 Save and close the file.

5 Enter the following command to restart the tomcat services. This enables tomcat to locate the certification files.

```
[root@vodwater opt]# service tomcat restart
```

6 Edit the **/etc/httpd/user-conf/ssl.ports** file and add the following line to the end of the file.

```
listen dncseth:443
```

7 Save and close the file.

8 Open the **/etc/httpd/user-conf/443.auth.conf** file in a text editor and complete the following steps to provide SSL access for the ECS and VCS Console.

```
[root@vodwater opt]# vi /etc/httpd/user-conf/443.auth.conf
```

   a Go to the **Uncomment to secure by host or subnet** section and uncomment out each entry from <Location /> to </Location>.

   b Open a line after Allow from dncseth and add an **Allow from [network IP prefix where the ECS and VCS Console resides]** entry. In this example, **Allow from 10.90.44** was added.

```
# Uncomment to secure by host or subnet
<Location />
 Order Allow,Deny
 Allow from 127.0.0.1
 Allow from dncs berlin
 Allow from appservatm
 Allow from dncseth
 Allow from 10.90.44

  ErrorDocument 403 "<html><head><title>Error 403</title></head><body><h2>SECURITY WARNING</h2>Web
 connections are only allowed from localhost.</body></html>"
 </Location>
```

   c Save and close the file.

9 On the EC or DTACS server, run the **gen_cert** command and select **Option 5** to check for dependencies and to enable apache SSL.

```
[root@vodwater opt]# /etc/httpd/gen_crt
```

10 Enter the following command to restart the **httpd** service.

```
[root@vodwater opt]# service httpd restart
[root@vodwater httpd]# service httpd-dncsws restart
```

11 Go to the ECS and the BOA nodes, respectively, and verify that the certificates successfully deployed.

Regionalizing the EC or DTACS Server to the ECS

Complete the following procedure to regionalize an SR 8.0 EC server or a DTACS 5.0 server to the ECS 3.0 system.

1 On the EC/DTACS server, enter the appropriate command to run the Add DNCS Feature script with the reset option.

EC:

```
[root@vodwater ~]# /dvs/dncs/etc/.ADF reset
```

DTACS:

```
[root@vodwater-DTACS50 ~]# /dvs/dtacs/etc/.ADF reset
```

2 Open the following file in a text editor to verify that the encryption keys for the server you want to register includes its unique encryption keys.

```
[root@vodwater ~]# vi /etc/consul/client.json
```

Example:

3 Does the file contain the appropriate entries for the EC or DTACS server?

- If **yes**, save and close the file.
- If **no**, either add the appropriate entries or refer to the Configuring RPC Encryption With TLS section in the ***SR 8.0 Installation and Migration Guide*** or the ***DTACS 5.0 Installation and Migration Guide***. Then save and close the file.

4 Enter the appropriate command to verify whether or not the EC or DTACS server is currently registered to an ECS.

EC:

```
[root@vodwater ~]# echo "select * from registration_config"|
dbaccess dncsdb
```

DTACS:

```
[root@vodwater-DTACS50 ~]# echo "select * from
registration_config"| dbaccess dtacsdb
```

5 Were any rows found in the database for regionalization?

- If **yes**, go to Step 6.
- If **no**, go to Step 7.

6 Enter the appropriate command to delete the regionalization entry from the database. Then go to Step 7.

EC:

```
[root@vodwater ~]# echo "delete from
registration_config"|dbaccess dncsdb
```

DTACS:

```
[root@vodwater-DTACS50 ~]# echo "delete from
registration_config"|dbaccess dtacsdb
```

7 Open the following file in a text editor.

```
[root@vodwater ~]# vi
/etc/httpd/user-conf/CSCOec.bossreq.auth.conf
```

8 Go to the second occurrence of "SUCCESS" and open a line above it.

9 Add the following line, substituting the IP address of the BOA node for [BOA_IP_ADDRESS].

**%{REMOTE_ADDR} eq "[BOA_IP_ADDRESS]" or **

Example:

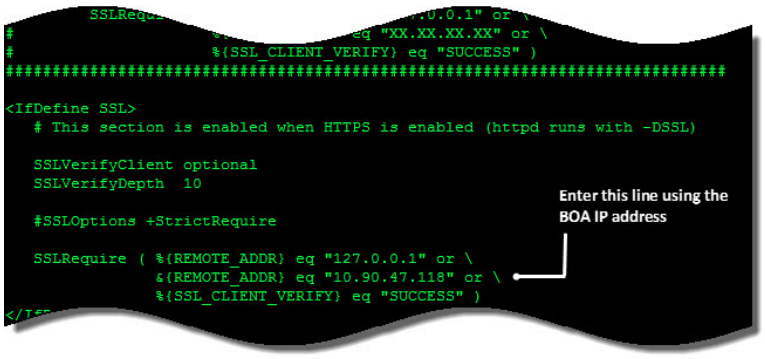

10 From a supported Web browser, login to the ECS Web UI.

11 Click the **Navigation** button, . and select **Console Admin > Host / IP White List**.

12 Click the + icon. The Add Row window appears.

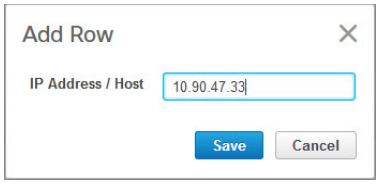

13 Enter the the IP Address for the EC or DTACS server you want to regionalize.

14 Click **Save**. The IP entry is added to the Host / IP White List.

15 Enter the following command to tail the eventManager log and monitor the registration process.

EC:

```
[root@vodwater ~]# tail -f /dvs/dncs/tmp/eventManager.000
```

DTACS:

[root@vodwater-DTACS ~]# `tail -f /dvs/dtacs/tmp/eventManager.000`

16 Access the EC or DTACS Web UI and select **EC > System Provisioning > Regionalization Configuration**.

17 Refer to one of the following tables to configure regionalization for an EC or a DTACS server.

- **EC Regionalization Parameters**

Field	Definition
EC ID	Hostname for the primary EC server; populated by default **Note**: If the primary and secondary hostnames are not unique, you can modify the hostname.
Primary EC Description	Text to describe the primary EC
EC Management URL Scheme	Maintain default
EC Management URL	Maintain default
Standby EC Description	Text to describe the secondary EC
Standby EC Timezone	Select the appropriate timezone from the dropdown menu
Standby EC IP	Enter the IP address for the secondary EC server
BOA URL	(Optional) Enter the destination host (and optional port), as well as the URL string to proxy BOSS requests to BOA

Example for EC:

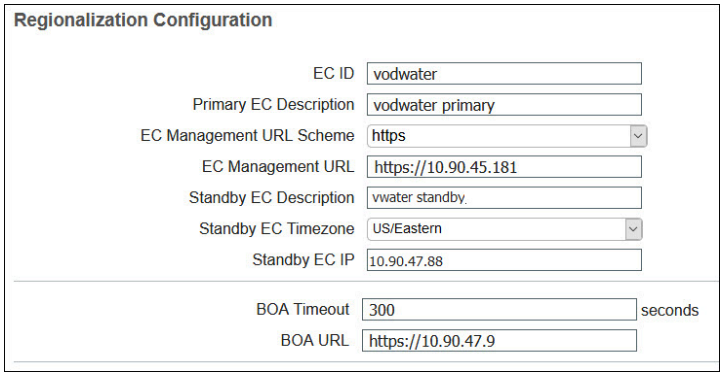

- **DTACS Regionalization Parameters**

Field	Definition
DTACSID	Hostname for the primary DTACS server; populated by default **Note**: If the primary and secondary hostnames are not unique, you can modify the hostname.
DTACS Description	Text to describe the primary DTACS
DTACS Management URL Scheme	Maintain default
DTACS Management URL	Maintain default
Standby DTACS Description	Text to describe the secondary DTACS
Standby DTACS Timezone	Select the appropriate timezone from the dropdown menu
Standby DTACS IP	Enter the IP address for the secondary DTACS server

Example for DTACS

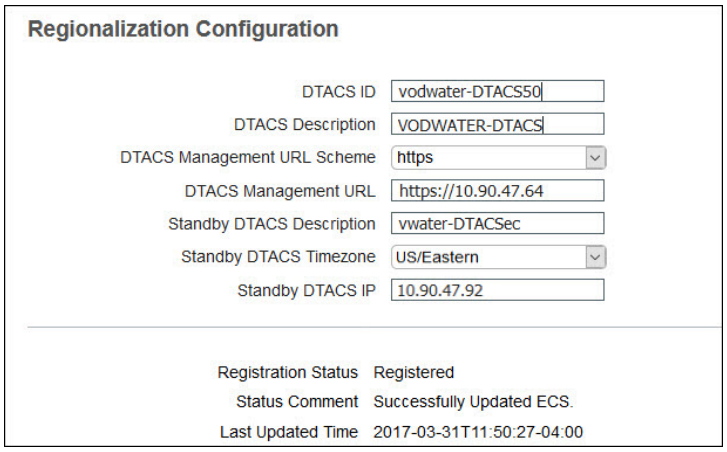

18 Click **Save**. The registration status will update to **Registered**.

19 Enter the appropriate command to verify that the registration status in the database is set to **Successfully Updated ECS**.

Command on EC:

```
[root@vodwater ~]# echo "select * from registration_config"|
dbaccess dncsdb
```

Command on DTACS:

```
[root@vodwater-DTACS50 ~]# echo "select * from
registration_config"| dbaccess dtacsdb
```

Result: The following entry should be present in the output.

```
status_comment      Successfully Updated ECS.
```

20 Log into the ECS Web UI and click the **Navigation** button, , and select **Control Plane > ECS Dashboard**.

21 Drill down in the ECS directory tree to verify the following:

- The EC or DTACS server you registered is present in the list
- A green checkmark, , is present on the icon for that server.

22 From the **Service Instance** List, you should have five new instances as shown below.

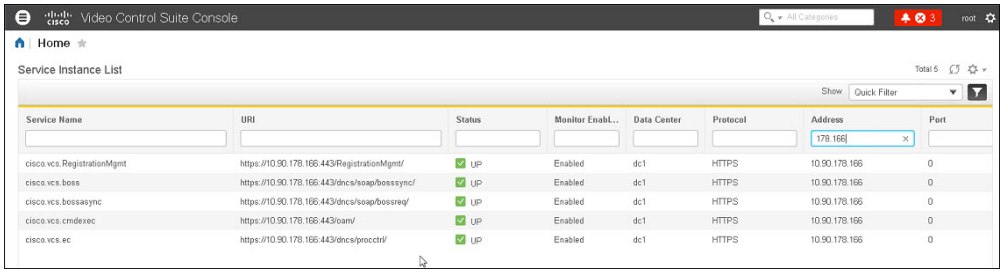

23 Click the **Network Element Management** tab and then click **Network Element Access Management**. The Network Element Access window appears.

24 To allow an ECS user to access the network element (i.e. EC, DTACS) you just registered, select the user who needs access from the **User** dropdown menu.

25 From the **Network Element** area, select that user(s) you want to authorize to access the server.

Note: Use the Shift key to select a sequence of users or the Ctrl key to select various users.

26 Click **Save**. A message appears in the lower right of the screen with the results of the change.

Note: It might take several minutes for the user to become authorized.

Verifying SNMP Configuration

After a successful EC registration, SNMP should be ready to forward alarms and alerts to the ECS. However, it may be possible that the SNMP community was deleted.

Complete the following steps to ensure that SNMP is still configured.

1. Using an SSH client, log into the EC that was just registered to the ECS.
2. Change to **root** user.

   ```
   [admin@vodwater ~]$ sudo -i
   ```

3. Enter the following command to verify that the snmpTrapHandler log level is set to debug (+DE).

   ```
   [root@vodwater ~]# logLvl| grep -i snmptrap
   ```

4. Did the output show +DE?
   - If **yes**, go to Step 5.
   - If **no**, enter the following command and then go to Step 5.

     ```
     [root@vodwater ~]# logLvl snmpTrapHandler +DE
     ```

5. Enter the following command to open the **snmpTrapHandler** log file.

   ```
   [dncs@vodwater ~]# less /dvs/dncs/tmp/snmpTrapHandler.000
   ```

6. Scroll through the file to look for an alarm or alert.

   

7. Are events or alarms present in the file?
   - If **yes**, go to Step 24.
   - If **no**, go to the next step.
8. Change to **dncs** user.

   ```
   [root@vodwater ~]# su dncs
   ```

9 Enter the following command to verify that the SNMP community exists.

```
[dncs@vodwater root]$ config_snmp_users
```

10 Is the **ciscointernalcomm** community present?

- If **yes**, enter **q** and then go to Step 24.
- If **no**, enter **a** to add a new user/community.

11 When prompted for the User/Community, type **ciscointernalcomm** and press **Enter**.

12 When prompted for the SNMP version, enter **2** and press **Enter**.

13 Type **l** (lower case L) to view the list of users/communities. The new entry is added to a **Pending** users section.

14 Enter **q** to quit. A message to bounce the snmpTrapHandler process appears.

15 Log into the EC Web UI.

16 Select the **snmpTrapHandler** process and click **Stop**. The service will change to Stopped.

17 Select the **snmpTrapHandler** process and click **Start**. The service will change to Active.

18 Repeat Step 9 to verify that the new user/community is now listed in the **Current users** section.

19 Type **exit** to change back to root user.

20 Type the following command to verify that the ciscointernalcomm community was added to the **snmpd.conf** file.

```
[root@vodwater ~]# grep ciscointernalcomm
/dvs/dncs/etc/snmpd.conf
```

Output:

```
trapsess -v 2c -c ciscointernalcomm localhost:162
```

21 Is the **trapsess** line present?

- If **no**, go to Step 22.
- If **yes**, go to Step 24.

22 Open the **snmpd.conf** file in a text editor and add the trapsess line to the end of the file.

```
[root@vodwater ~]# vi /dvs/dncs/etc/snmpd.conf
```

23 Save and close the file.

24 Enter the following command to verify that the ownership of the snmpd.conf file is **dncs:dncs**.

```
[root@vodwater ~]# ls -ltr /dvs/dncs/etc/snmpd.conf
```

25 Is the ownership of the file **dncs:dncs**?

- If **no**, enter the following command to change the ownership to dncs:dncs. Then go to Step 26.

  ```
  [root@vodwater ~]# chown dncs:dncs /dvs/dncs/etc/snmpd.conf
  ```

- If **yes**, go to Step 28.

26 Enter the following command to restart the **snmpd** service.

```
[root@vodwater ~]# service snmpd restart
```

27 Repeat Step 20 to verify that the **ciscointernalcomm** community was added to the **snmpd.conf** file.

28 Enter the following command to verify that you can see events/alarms in the **snmpTrapHandler** log file.

```
[root@vodwater ~]# less /dvs/dncs/tmp/snmpTrapHandler.000
```

Verifying ECS Functionality After Regionalizing a Client

Once you have successfully registered a network element to the ECS, you should verify the ECS functionality. The following features should be tested from the ECS Web UI.

Note: When executing these tests, go to the **/opt/jboss-as/standalone/log** directory to monitor the appropriate log.

- **Operations and Administration (OAM)**
  - **Command Execution**: execute system commands on the remote EC/DTACS system
  - **Backup and Restore**: verify a software backup
- **Regional Provisioning (RPS)**
  - **Export from EC Device**: test that you can export data from the EC database and generate the export data file
  - **Import XML Template**: generate an apply (ingest) file from an XLS template based on the selected provisioning element (PE), and load the generated ingest file
  - **Apply to EC Device(s)**: apply (ingest) a request from the RPS to the RPO (Regional Provisioning Orchestrator), a service running on the EC. The RPO should process the apply request and then fetch the apply file using SFTP. It should then parse the file and provisions the data into the EC
- **CPE Management**
  - **Manage CPE**: verify that you can query a CPE
  - **Batch Management**: test that you can perform a batch install

Import CPE Data From the EC to the ECS

Important:

- This procedure is written for systems that are built in an isolated network environment and addresses the DE21527 defect.
- This procedure is only executed on an EC server.

Complete the following procedure to export CPE data from the EC and then export it to the ECS.

1 From the EC window, change to **root** user.

```
[admin@vodwater ~] sudo -i
```

2 Create a directory in **/dvs/backups/dncs**.

Note: For this example, we will create the tmp directory.

```
[root@vodwater ~]# mkdir /dvs/backups/tmp
```

3 Change to the **/dvs/backups/tmp** directory.

```
[root@vodwater ~]# cd /dvs/backups/tmp
```

4 Execute the following command to export the MAC addresses from the EC and save them to a file on the EC.

Note: In this example, we are using macaddress as the file name. You can use any file name of your choice.

```
[root@vodwater ~]# nohup java -jar /dvs/dncs/bin/ExportCPE.jar
-outfile macaddress -thread 64
```

Result: A zip file is created for the list of MAC address. In this example, the file is named macaddress.zip.

5 Copy the **ImportCPE** script to the current directory.

```
[root@vodwater ~]# cp -p /dvs/dncs/bin/ImportCPE.jar .
```

6 Execute the following command to change access permissions to the directory.

```
[root@vodwater ~]# chmod -R 755 .
```

7 As **admin** user on the ECS node, change to the **/home/admin** directory.

8 Using scp, copy the following two files from the EC to the **/home/admin** directory on the ECS node.

Command Syntax:

```
scp [userID]@EC-IP:/[directory]/[filename] .
```

- The zip file created in Step 4 (e.g. macaddress.zip)

Example:

```
[admin@ECS admin]# scp
admin@10.90.47.18:/dvs/backups/tmp/macaddress.zip .
```

- The /dvs/backups/tmp/ImportCPE.jar file

Example:

```
[admin@ECS admin]$ scp
admin@10.90.47.18:/dvs/backups/tmp/ImportCPE.jar .
```

9 Execute the following command to import the CPEs to the ECS.

Notes:

- Replace [filename.zip] with the name of the file you generated in Step 4.
- Replace [RAC_Scan] with the RAC Scan hostname or IP address for your ECS 3.0 system.
- Replace [ECS_Username] with the username for your ECS database user.
- Replace [ECS_Password] with the password for your ECS database user.
- If the password contains special characters, they will need to be "escaped" (i.e. -pwd "2g3n3r/!c").

Command Syntax:

```
nohup java -jar ImportCPE.jar -file [filenames.zip] -jdbc
jdbc:oracle:thin:@[RAC_Scan]:1535/CABHE -uid [ECS_Username]
-pwd [ECS_Password]
```

Example:

```
[admin@ECS admin]$ nohup java -jar ImportCPE.jar -file
macaddress.zip -jdbc jdbc:oracle:thin:@172.20.36.5:1535/CABHE
-uid NEXTX_ECS -pwd nextxtest
```

10 Repeat this procedure if you have other ECs registered to an ECS.

Additional Features for Regionalization

Deleting a Registration from an ECS

1. Enter the following command to tail the eventManager log and monitor the registration deletion process.

   EC:

   ```
   [root@vodwater ~]# tail -f /dvs/dncs/tmp/eventManager.000
   ```

   DTACS:

   [root@vodwater-DTACS ~]# `tail -f /dvs/dtacs/tmp/eventManager.000`

2. From the **Service Provisioning > Regionalization Configuration** Web UI, click **Unregister**.
3. When prompted to enter the Web UI credentials for the EC or DTACS server, enter your administrative username and password. Then click **Log in**. You are directed to the EC WUI where a deleted successfully message appears in the lower right corner of the window.

   Note: The status is updated to UnRegistered.

Registration Status	UnRegistered
Status Comment	Unregistered Successfully
Last Updated Time	2017-03-31T12:11:37-04:00

4. Enter the appropriate command to verify that the registration status in the database is now **Unregistered Successfully**.

   Command for EC:

   echo "select * from registration_config" | dbaccess dncsdb -

   Command for DTACS:

   echo "select * from registration_config" | dbaccess dtacsdb -

   Example Output:

   ```
   status_comment      Unregistered Successfully
   ```

5. From the Regionalization Web UI, click **Delete** to delete the registration to the ECS.
6. When prompted to confirm the deletion, click **OK**.
7. Enter the appropriate command from Step 3 to verify that the registration is removed from the database.

   Example Output:

   ```
   Database selected.
   No rows found.
   Database closed.
   ```

8. Verify that the **rpa.config** file was deleted from the /dvs/dvsFiles/rpa directory:

   ```
   ls -ltr /dvs/dvsFiles/rpa
   ```

9 Log into the ECS Web UI and click the **Navigation** icon, , and select **Control Plane > ECS Dashboard**.

10 Drill down in the ECS directory tree to verify that the EC or DTACS server you deleted is no longer in the list.

11 From the Service Instance List, verify that the five instances associated with this network element are no longer present.

Registering an EC or a DTACS Server to a Different ECS

1 Complete the procedure in the following sections:

- ***Deleting a Registration from an ECS*** (on page 183)
- ***Regionalizing the EC or DTACS Server to the ECS*** (on page 172)
- ***Verifying SNMP Configuration*** (on page 177)

2 Enter the appropriate command to verify that the database indicates that the EC or DTACS server is now registered to the new ECS.

Command on EC:

```
[root@vodwater ~]# echo "select * from registration_config"|
dbaccess dncsdb
```

Command on DTACS:

```
[root@vodwater-DTACS50 ~]# echo "select * from
registration_config"| dbaccess dtacsdb
```

Result: The following entry should be present in the output.

```
status_comment      Successfully Updated ECS.
```

3 Log into the ECS Web UI and click the **Navigation** icon and select **Control Plane > ECS Dashboard**.

4 Drill down in the ECS directory tree to verify the following:

- The EC or DTACS server you registered is present in the list
- A green checkmark, , is present on the icon for that server.

5 From the Service Instance List, you should have five new instances listed.

6 Complete the steps in ***Verifying ECS Functionality*** (on page 140).

Retrying a Registration Request

There may be instances of a temporary network outage or timeout that results in a "timedout" registration status. If this occurs and you do not need to redefine any values in the Regionalization Configuration Web UI, complete these steps to retry the registration.

1 In the EC Regionalization Configuration WUI, click **Retry**. The registration status goes from **Pending** to **Registered**.

2 Enter the appropriate command to verify that the database indicates that the EC or DTACS server is now registered to the ECS.

Command on EC:

```
[root@vodwater ~]# echo "select * from registration_config"|
dbaccess dncsdb
```

Command on DTACS:

```
[root@vodwater-DTACS50 ~]# echo "select * from
registration_config"| dbaccess dtacsdb
```

Result: The following entry should be present in the output.

```
status_comment      Successfully Updated ECS.
```

3 Log into the ECS Web UI and click the **Navigation** icon and select **Control Plane > ECS Dashboard**.

4 Drill down in the ECS directory tree to verify the following:

- The EC or DTACS server you registered is present in the list
- A green checkmark,  , is present on the icon for that server.

5 From the Service Instance List, you should have five new instances present.

C

Configure Local Sign On

This appendix includes the procedure to manually configure local sign on when operating in regional mode. **Manually disabling SSO (single sign on) should only be a temporary as the features for RPS will not be functional.**

In This Appendix

Enabling Local Sign On

Complete the following procedure to enable local sign on when your system is in a regionalized mode.

From the terminal window, log into the EC or DTACS system where you want to enable local sign on as **admin** user.

1 As **root** user, enter the following command to change to the **/etc/httpd/conf.cisco** directory.

```
[root@EC/DTACS ~]# cd /etc/httpd/conf.cisco
```

2 Enter the following command to verify that the CSCOec.auth file is pointing to **CSCOec.auth.sso**.

```
[root@EC/DTACS conf.cisco]# ls -ltr
```

```
-r--r-----. 1 root root  471 Apr 12 13:28 CSCOrptwrt_httpd.conf
-r--r-----. 1 dncs dncs  357 May 16 12:11 CSCOapsrv.conf
-r--r-----. 1 dncs dncs 1959 May 17 06:33 CSCOec-webui.worker1.conf
-r--r-----. 1 dncs dncs 2460 May 17 06:33 CSCOec-webui.gen.conf.orig
-r--r-----. 1 dncs dncs 2460 May 17 06:33 CSCOec-webui.gen.conf
-r--r--r--. 1 root root  736 May 17 06:33 CSCOec.soap.conf.orig
-rw-r--r--. 1 root root  736 May 17 06:33 CSCOec.soap.conf
-r--r--r--. 1 root root 1663 May 17 06:33 CSCOec.conf.orig
-rw-r--r--. 1 root root 1663 May 17 06:33 CSCOec.conf
-rw-r--r--. 1 root root 1037 May 17 06:33 80.http.vhost.orig
-rw-r--r--. 1 root root 1037 May 17 06:33 80.http.vhost
-rw-r--r--. 1 root root 3652 May 17 06:33 443.https.vhost.orig
-rw-r--r--. 1 root root 3652 May 17 06:33 443.https.vhost
-rw-r--r--. 1 root root  584 May 18 11:40 CSCOoam.conf
-rwxr-xr-x. 1 dncs dncs  350 May 18 11:45 CSCORegistrationMgmt.conf
-r--r-----. 1 dncs dncs  364 May 22 14:30 CSCOneDataAccess.conf
lrwxrwxrwx. 1 root root   39 May 23 19:38 CSCOec.auth.presso -> /etc/httpd/user-conf/CSCOec.auth_digest
lrwxrwxrwx. 1 root root   36 May 23 19:39 CSCOec.auth -> /etc/httpd/user-conf/CSCOec.auth_sso
```

Indicates SSO is enabled

3 Enter the following command to go to the **/etc/httpd/rpa** directory.

```
[root@EC/DTACS conf.cisco]# cd ../rpa
```

4 Enter the following command to launch the SSO menu.

```
[root@EC/DTACS rpa]# ./set_sso
```

```
Choose from following options:

1. Enable SSO

2. Disable SSO

3. Show SSO status

4. Exit

Please enter your choice: [1|2|3|4]
```

5 Enter **2** to disable SSO. SSO is disabled and the http processes are restarted.

Note: Disabling SSO will enable local sign on.

```
Please enter your choice: [1|2|3|4]
2
Disabling SSO
Restoring prior authentication settings, see CSCOec.auth
Restarting httpd
Stopping httpd:                                            [  OK  ]
Starting httpd:                                            [  OK  ]
Restarting httpd-dncsws
Stopping httpd-dncsws:                                     [  OK  ]
Starting httpd-dncsws:                                     [  OK  ]
```

6 Enter the following command to verify that CSCOec.auth is now pointing to **CSCOec.auth.digest**.

```
[root@EC/DTACS rpa]# ls -ltr ../conf.cisco
```

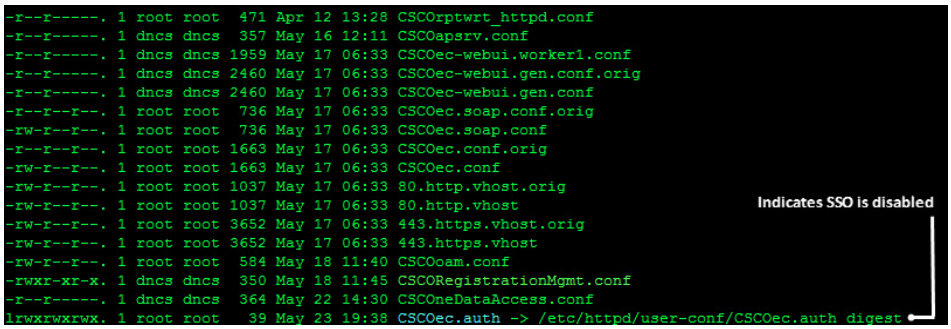

7 From a supported Firefox Web browser, log into the EC/DTACS Web UI and verify that you can now log in with your local credentials.

Disabling Local Sign On

Complete the following procedure to disable local sign and enable SSO.

1 From the terminal window, log into the EC or DTACS system where you want to disable local sign on as **admin** user.

2 As **root** user, enter the following command to change to the **/etc/httpd/conf.cisco** directory.

```
[root@EC/DTACS ~]# cd /etc/httpd/conf.cisco
```

3 Enter the following command to verify that the CSCOec.auth file is pointing to **CSCOec.auth.digest**.

```
[root@EC/DTACS conf.cisco]# ls -ltr
```

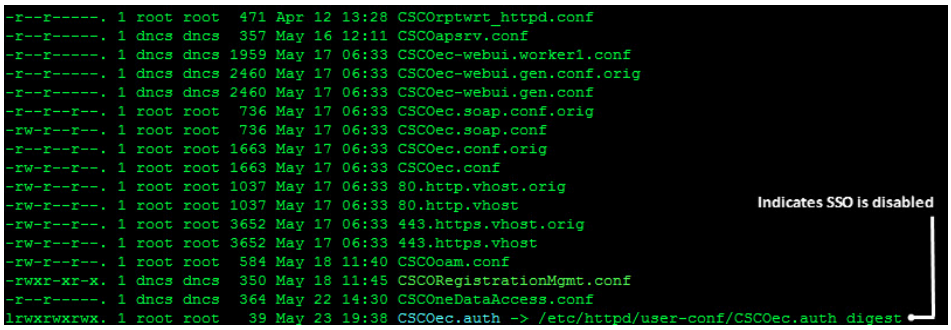

4 Enter the following command to go to the **/etc/httpd/rpa** directory.

```
[root@EC/DTACS conf.cisco]# cd ../rpa
```

5 Enter the following command to launch the SSO menu.

```
[root@EC/DTACS rpa]# ./set_sso
```

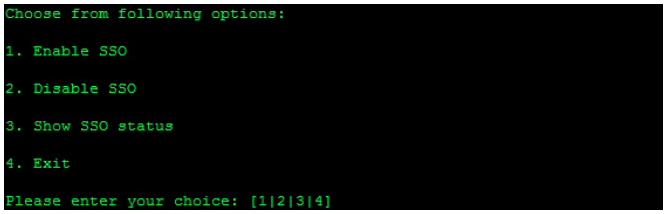

6 Enter **1** to enable SSO. SSO is enabled and the http processes are restarted.

Note: Enabling SSO will disable local sign on.

7 Enter the following command to verify that CSCOec.auth is now pointing to **CSCOec.auth.sso**.

```
[root@EC/DTACS rpa]# ls -ltr ../conf.cisco
```

8 From a supported Firefox Web browser, log into the EC/DTACS Web UI and verify that you can now log in with your SSO credentials.

D

ECS 3.0 Upgrade

This chapter provides the procedures to upgrade your existing VCS Console, ECS and/or BOA nodes to a new version of software.

Important: The RPM files for all new packages should have already been downloaded to the Admin Node and deployed to your NextX repo. If they have not, refer to the **Admin Node Installation Guide** for details.

In This Appendix

Preparing the Primary and Secondary VMs for Cloning

Important: Complete this procedure for an upgrade on the VCS Consoles, ECSs, or BOA VMs.

Complete the following procedures to prepare the primary and secondary VMs for cloning.

1. In unique terminal windows, log into the *primary* and *secondary* VM as **admin** user.
2. On the *primary* VM, enter the following command.

   ```
   [admin@platform ~]$ sudo sed -i "s/SUBSYSTEM.*//"
   /etc/udev/rules.d/70-persistent-net.rules
   ```
3. Repeat Step 2 on the *secondary* VM.
4. From the vSphere Web UI, select and right-click the *primary* VM and then select **Shut Down Guest OS**.
5. When prompted to confirm the shutdown for each VM, click **Yes**.
6. Repeat Steps 4 through 5 on the *secondary* VM.
7. Right-click the *primary* VM and select **Rename**.
8. In the **Enter a new name** text box, type a name that will reflect that it is the cloned VM.

   Note: We suggest appending verbage that indicates it is the cloned VM. In our example, we will append "-clone".
9. Repeat Step 7 through 8 on the *secondary* VM.

Cloning the Primary and Secondary VMs

Important: Complete this procedure for an upgrade on the VCS Consoles, ECSs, or BOA VMs.

Complete the following procedure to clone the primary and secondary VMs.

1. From the vCenter Web UI, select and right-click the *primary* VM and select **Clone to Virtual Machine**. The Clone Existing Virtual Machine window appears.
2. From the **Enter a name for the virtual machine** text box, enter the original name of the VM.
3. Click **Next**.
4. From the **Select a compute resource** window, select the same host where the original VM resides.
5. View the **Compatibility** area in the window and verify that the **Validation succeed** message appears.
6. Click **Next**.
7. From the **Select storage** window, maintain the default for the virtual disk format and then select the same datastore that the original VM is using.
8. Click **Next**.
9. Click **Next** again.
10. Review the settings for the new VM and then click **Finish**. The cloning process starts.
11. From the **Recent Tasks** area, monitor the creation of the cloned VM until the task completes.
12. Repeat Steps 1 through 11 for the *secondary* VM.

Upgrading the VCS Console Servers

Complete the following steps to upgrade the *original* VCS Console servers.

Important:

- The new vcsconsole RPM should already be present in the NextX repo. If it is not, refer to the **Updating the Application Packages Repo** section in the ***Admin Node Installation Guide*** to update the NextX repo with the new software.
- Do not execute a "yum update" command on the VCS Console as the vcsconsole service does not support upgrading using this method.

1 From the vSphere Web UI, right-click the *original primary* server and select **Power On**.

2 Verify basic functionality of the server.

3 Enter the following command to create a backup directory.

Command Syntax:

```
sudo mkdir ~admin/bkp-nextx_[date]
```

Example:

```
[admin@vcsconsole ~]$ sudo mkdir ~admin/bkp-nextx_20170421
```

4 Enter the following command to backup the **server.xml** configuration file to the new backup directory.

```
[admin@vcsconsole ~]$ sudo cp -p
/opt/web/vcsconsole/conf/server.xml
~admin/bkp-nextx_20170421/server.xml.vcs
```

5 Enter the following command to stop the **vcsconsole** service.

```
[admin@vcsconsole ~]$ sudo service vcsconsole stop
```

6 Repeat Steps 1 through 5 on the *original secondary* VCS Console.

7 On the *primary* server, enter the following command to install the new VCS Console RPM package.

Command Syntax:

```
sudo yum install vcsconsole-[VERSION]
```

Example:

```
[admin@vcsconsole ~]$ sudo yum install vcsconsole-4.0.19
```

8 Complete the following steps ***only*** on the *primary* VCS Console.

a Enter the following command to update the VCS Console schema.

```
[admin@vcsconsole ~]$ sudo /opt/vcs/bin/vcsutils.sh
-updateSchema
```

b Enter the following command to populate the VCS Console database.

```
[admin@vcsconsole ~]$ sudo /opt/vcs/bin/vcsutils.sh
-populateDB
```

9 On the *secondary* server, repeat Step 7 to install the new VCS Console package.

10 In the terminal window on the *primary* VCS Console, enter the following command to restore the **server.xml.vcs** file that you backed up in Step 4.

```
[admin@vcsconsole ~]$ sudo cp
~admin/bkp-nextx_20170421/server.xml.vcs
/opt/web/vcsconsole/conf/server.xml
```

11 Enter the following command to stop the **Alarm_Manager** service.

```
[admin@vcsconsole ~]$ sudo service Alarm_Manager stop
```

12 Enter the following command to reboot the server.

```
[admin@vcsconsole ~]$ sudo reboot
```

13 Log back into the server as **admin** user.

14 Enter the following commands to verify that the vcsconsole, Alarm_Manager and DECAP services are all running.

```
[admin@vcsconsole ~]$ sudo service vcsconsole status
[admin@vcsconsole ~]$ sudo service Alarm_Manager status
```

```
[admin@vcsconole ~]$ sudo service vcsconsole status
nds_vcsconsole process(es) running (30851)
[admin@vcsconsole ~]$ sudo service Alarm_Manager status
Alarm_Manager process(es) running (15292)
Decap process(es) running (15364 15369 )
```

15 Are these services running?

- If **yes**, go to the next step.
- If **no**, type the appropriate command to start the service.

  Note: If either or both services fail to start, contact Cisco Services.

  ```
  [admin@vcsconsole ~]$ sudo service vcsconsole start
  [admin@vcsconsole ~]$ sudo service Alarm_Manager start
  ```

16 Repeat Steps 10 through 15 on the *secondary* server.

17 From a Web browser, verify that you can access the VCS Console using the VIP.

URL Format:

```
https://[VCS Console VIP]:6605/vcsconsole
```

Example:

```
https//10.90.47.33:6605/vcsconsole
```

18 Verify functionality for both the *primary* and *secondary* VCS Consoles.

19 Is there an update to **Alarm_Manager**?

- If **no**, you have completed this procedure.
- If **yes**, go to the next step.

20 Enter the following command to stop the **Alarm_Manager** service.

```
[admin@vcsconsole ~]$ sudo service Alarm_Manager stop
```

21 Enter the following command to update **Alarm_Manager**.

```
[admin@vcsconsole ~]$ sudo yum update Alarm_Manager
```

22 When prompted to confirm the installation, type **y** and press **Enter**.

23 Enter the following command to start **Alarm_Manager**.

```
[admin@vcsconsole ~]$ sudo service Alarm_Manager start
```

24 Enter the following command to verify that the **Alarm_Manager** and the **Decap** processes started successfully.

```
[admin@vcsconsole ~]$ service Alarm_Manager status
```

Result: An Alarm Manager and a Decap process is running message appears along with each respective PID number.

Example:

```
[admin@vcsconsole ~]$ sudo service Alarm_Manager status
Alarm_Manager process(es) running (15292)
Decap process(es) running (15364 15369 )
```

25 Enter the following command to review the **catalina.out** file and confirm that there are no warnings or errors in relation to the Alarm_Manager service.

```
[admin@vcsconsole ~]$ less
/opt/apache/tomcat70/logs/catalina.out
```

26 Enter the following command to install the **Alarm Manager Web UI**.

```
[admin@vcsconsole ~]$ yum update AlarmManagement_UI
```

27 When prompted, enter **y** to confirm the installation.

28 Repeat Steps 19 through 27 on the *secondary* VCS Console.

Upgrading the ECS Servers

Complete the following procedure to upgrade the primary and secondary ECS servers in your system.

Important: The new ECS RPM should already be present in the NextX repo. If it is not, refer to the **Updating the Application Packages Repo** section in the ***Admin Node Installation Guide*** to update the NextX repo with the new software.

1. From the vSphere Web UI, right-click the *original primary* server and select **Power On**.
2. Verify basic functionality of the server.
3. Enter the following command to upgrade the ECS application.

   ```
   [admin@ecs ~]$ sudo yum update
   ```
4. When prompted to confirm the updates, type **y** and press **Enter**.
5. When the updates complete, enter the following command to reboot the ECS node.

   ```
   [admin@ecs ~]$ sudo reboot
   ```
6. Log back into the ECS as **admin** user.
7. Verify functionality.
8. Repeat Steps 1 through 7 on the *original secondary* ECS node.

Upgrading the BOA Servers

Complete the following procedure to upgrade the primary and secondary BOA servers in your system.

Important: The new billingAdaptor RPM should already be present in the NextX repo. If it is not, refer to the **Updating the Application Packages Repo** section in the ***Admin Node Installation Guide*** to update the NextX repo with the new software.

1. From the vSphere Web UI, right-click the *original primary* server and select **Power On**.
2. Verify basic functionality of the server.
3. Enter the following command to upgrade the BOA application and BOA UI, as needed.

   ```
   [admin@boa ~]$ sudo yum update
   ```
4. When prompted to confirm the updates, type **y** and press **Enter**.
5. When the updates complete, enter the following command to redeploy the Billing Adaptor.

   ```
   [admin@boa ~]$ sudo
   /opt/cisco/billingadaptor/bin/billingadaptor.sh redeploy
   ```

   Note: If the Billing Adaptor failed to start, refer to the **Installing the Billing Adaptor User Guide** to assist with troubleshooting.
6. Log into the VCS Console UI. The Service Instance List displays.
7. Verify that all services are present which a status of **UP**.
8. Repeat Steps 1 through 7 on the *original secondary* BOA node.

E

Patch Installs

This appendix describes the procedures to install a patch to a primary and secondary node in your NextX ECS system.

In This Appendix

Installing a Patch to the ECS Nodes

This section describes the procedures to install a patch to the primary and secondary ECS in your NextX system.

The format for an ECS patch is: **CSCOecs-patch-[VERSION].[DATE].[PLATFORM].rpm**

The first patch that is available is always has "-2" extension for the version, for example, CSCOecs-patch-3.0.16-**2**.[VERSION].[DATE].[PLATFORM]. The naming convention for any future patches will be incrementally versioned, for example, CSCOecs-patch-3.0.16-**3**.[VERSION].[DATE].[PLATFORM].

Each new patch is a cumulative patch and includes all of the patches previous to the current package version. For example, CSCOecs-patch-3.0.16-**5**.[VERSION].[DATE].[PLATFORM] would include patches for 3.0.16-**2** through 3.0.16-**5**.

Preparing for the Patch Install

Complete the following steps prior to executing the patch install.

1 Are you using vSphere Web UI or vSphere client?

- If **no**, refer to the ***Explorer Controller Suite 3.0 Backup and Restore User Guide*** to back up the Oracle database.
- If **yes**, clone your *primary* and your *secondary* ECS nodes to create a backup of the current ECS server.

Note: If necessary, refer to ***Preparing the Primary and Secondary VMs for Cloning*** (on page 194) and ***Cloning the Primary and Secondary VMs*** (on page 195) for details.

2 From a Web browser, enter the following command to verify that the patch has been deployed on the Admin Node that is associated with your system.

URL Syntax:

```
http://[Admin_Node_IP]/repos/nextx/8.0/
```

Example:

```
http://10.90.47.60/repos/nextx/8.0/
```

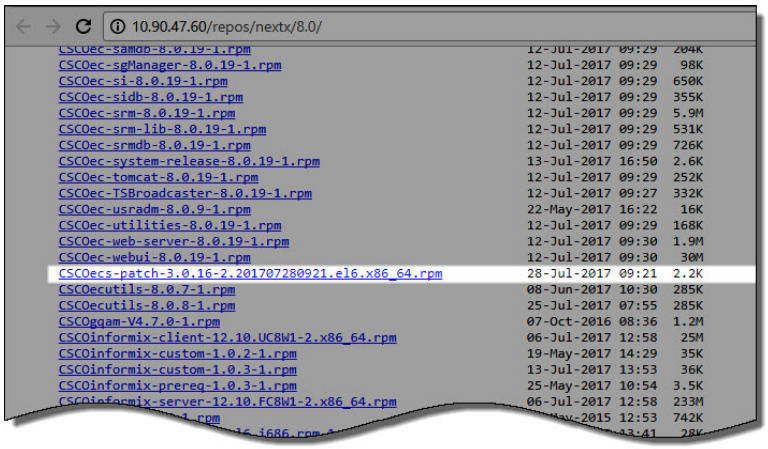

3 Is the **CSCOecs-patch** RPM present in the NextX repo?

- If **yes**, go to the next section.
- If **no**, refer to the **Updating the Application Packages Repo** in the ***Admin Node Installation Guide*** to update the NextX repo with the patch software.

Installing the Patch

Complete the following procedure to install the patch to the ECS nodes.

1 As **admin** user on the *primary* ECS node, enter the following command to verify if a CSCOecs-patch package is currently installed.

Note: If a CSCOecs-patch is not installed, this command will produce no output.

```
[admin@ecs ~]$ rpm -qa | grep -i CSCOecs-patch
```

2 Does the **CSCOecs-patch** package exist on the system?

- If **no**, enter the following command. The install process is set up and the dependencies are verified; then an **Is this ok [n/Y]** message displays.

  Note: You must enter the full package name in the "yum install" command if this is the first time installation of an ECS patch.

  Command Syntax:

  ```
  sudo yum install CSCOecs-patch-[VERSION].[DATE].[PLATFORM]
  ```

  Example:

  ```
  [admin@ecs ~]$ sudo yum install
  CSCOecs-patch-3.0.16-2.201707280921.el6.x86_64
  ```

- If **yes** and you are installing a newer CSCOecs-patch package, enter the following command. The install process is set up and the dependencies are verified; then an **Is this ok [n/Y]** message displays.

Note: A newer patch install does not require you to enter the full package name in the command.

```
[admin@ecs ~]$ sudo yum update CSCOecs-patch*
```

```
Loaded plugins: security
Setting up Install Process
Resolving Dependencies
--> Running transaction check
---> Package CSCOecs-patch.x86_64 0:3.0.16-2.201707280921.el6 will be installed
--> Processing Dependency: LoadPIMSService = 3.0.7-1 for package: CSCOecs-patch-3.0.16-2.201707
280921.el6.x86_64
--> Running transaction check
---> Package LoadPIMSService.noarch 0:3.0.6-1 will be updated
---> Package LoadPIMSService.noarch 0:3.0.7-1 will be an update
--> Finished Dependency Resolution

Dependencies Resolved

================================================================================================
 Package              Arch         Version                          Repository           Size
================================================================================================
Installing:
 CSCOecs-patch        x86_64       3.0.16-2.201707280921.el6        solution-base       2.2 k
Updating for dependencies:
 LoadPIMSService      noarch       3.0.7-1                          solution-base        29 M

Transaction Summary
================================================================================================
Install       1 Package(s)
Upgrade       1 Package(s)

Total download size: 29 M
Is this ok [y/N]:
```

3 Type **y** and press **Enter**. The installation continues and when finished, a **Complete!** message displays.

Note: In this particular patch install, the CSCOecs-patch and the LoadPIMSService packages were downloaded and installed. Subsequent patch installs will contain different RPMs as required.

```
Downloading Packages:
(1/2): CSCOecs-patch-3.0.16-2.201707280921.el6.x86_64.rpm               | 2.2 kB     00:00
(2/2): LoadPIMSService-3.0.7.rpm                                       |  29 MB     00:00
------------------------------------------------------------------------------------------
Total                                                   65 MB/s |  29 MB     00:00
Running rpm_check_debug
Running Transaction Test
Transaction Test Succeeded
Running Transaction
  Updating   : LoadPIMSService-3.0.7-1.noarch                                    1/3
  Installing : CSCOecs-patch-3.0.16-2.201707280921.el6.x86_64                    2/3
  Cleanup    : LoadPIMSService-3.0.6-1.noarch                                    3/3
  Verifying  : CSCOecs-patch-3.0.16-2.201707280921.el6.x86_64                    1/3
  Verifying  : LoadPIMSService-3.0.7-1.noarch                                    2/3
  Verifying  : LoadPIMSService-3.0.6-1.noarch                                    3/3

Installed:
  CSCOecs-patch.x86_64 0:3.0.16-2.201707280921.el6

Dependency Updated:
  LoadPIMSService.noarch 0:3.0.7-1

Complete!
```

4 Enter the following command to verify that the **CSCOecs-patch** package successfully installed, as well as any other packages.

Command Syntax:

```
rpm -qa | egrep "[package_name_1] | [package_name_2]
[package_name_n]"
```

Example:

```
[admin@ecs ~]$ rpm -qa | egrep -i "CSCOecs-patch|LoadPIMS"
```

```
CSCOecs-patch-3.0.16-2.201707280921.el6.x86_64
LoadPIMSService-3.0.7-1.noarch
```

5 Enter the following command to query the patch package and view the release date, the version, other installed packages, and the issues corrected in the patch.

Command Syntax:

```
rpm -q --changelog CSCOecs-patch-[VERSION].[DATE].[PLATFORM]
```

Example:

```
[admin@ecs ~]$ rpm -q --changelog
CSCOecs-patch-3.0.16-2.201707280921.el6.x86_64
```

```
* Wed Jul 26 2017 - 3.0.16-2
LoadPIMSService-3.0.7-1.rpm

- CSCvf36211 Batch installs of EMM tar files with an underscore in the name fails
```

6 Enter the following command to restart **jboss-as**.

```
[admin@ecs ~]$ sudo service jboss-as restart
```

7 Verify ECS server functionality.

8 Is your server functioning properly?

- If **yes** and you successfully installed the patch on the *primary* server, go to the next step.
- If **yes** and you successfully installed the patch on the *primary* and the *secondary* ECS servers, you have completed this procedure.
- If **no**, troubleshoot the system. If you cannot remedy the issue, contact Cisco Services.

  Note: You can also choose to uninstall the patch. Refer to the next section for details.

9 Repeat Steps 1 through 8 on the *secondary* ECS server.

Uninstalling the Patch on the ECS Server

Complete the following steps to uninstall the patch on the ECS node. This procedure also downgrades any packages that were installed/upgraded as dependencies to the patch installation.

1 Enter the following command to verify the current **CSCOecs-patch** version that is installed on the ECS node.

```
[admin@ecs ~]$ rpm -qa | grep -i CSCOecs-patch
```

```
CSCOecs-patch-3.0.16-2.201707280921.el6.x86_64
```

2 As **admin** user, enter the following command to obtain the **ID** of the CSCOecs-patch installation.

```
[admin@ecs ~]$ sudo yum history package-list CSCOecs-patch\*
```

```
Loaded plugins: security
ID     | Action(s)      | Package
-------------------------------------------------------------------------------
    43 | Install        | CSCOecs-patch-3.0.16-2.201707280921.el6.x86_64     EE
history package-list
```

3 From the **ID** column, record the ID number for the **CSCOecs-patch** installation.

ID Number: _____

4 Enter the following command to uninstall the patch using the ID number you recorded in the previous step. An **Is this ok [n/Y]** message displays.

Command Syntax:

```
sudo yum history undo [ID_number]
```

Example:

```
[admin@ecs ~]$ sudo yum history undo 43
```

```
Resolving Dependencies
--> Running transaction check
---> Package CSCOecs-patch.x86_64 0:3.0.16-2.201707280921.el6 will be erased
---> Package LoadPIMSService.noarch 0:3.0.6-1 will be a downgrade
---> Package LoadPIMSService.noarch 0:3.0.7-1 will be erased
--> Finished Dependency Resolution

Dependencies Resolved

================================================================================
 Package              Arch         Version                      Repository          Size
================================================================================
Removing:
 CSCOecs-patch        x86_64       3.0.16-2.201707280921.el6    @solution-base      0.0
Downgrading:
 LoadPIMSService      noarch       3.0.6-1                      solution-base        29 M

Transaction Summary
================================================================================
Remove        1 Package(s)
Downgrade     1 Package(s)

Total download size: 29 M
Is this ok [y/N]:
```

5 Type **y** and press **Enter**. The downgrade proceeds, and when finished, a **Complete!** message displays.

Note: In this example, the CSCOecs-patch is removed and the LoadPIMSService package is downgraded. Subsequent patches will have different dependencies as required.

6 Enter the following command to verify the current versions of the patches and any other packages that were downgraded.

Note: If this was the first time a patch was installed on the ECS, no output displays for the CSCOecs-patch package.

Command Syntax:

```
rpm –qa | egrep "[package_name_1] | [package_name_2] [package_name_n]"
```

Example:

```
[admin@ecs ~]$ rpm -qa | egrep -i "CSCOecs-patch|LoadPIMS"
```

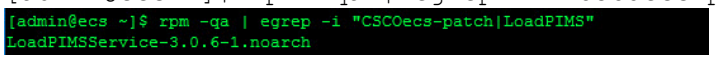

Enter the following command to restart **jboss-as**.

```
[admin@vodwater ~]$ sudo service jboss-as restart
```

7 Verify the functionality of the ECS server.

8 Is your server functioning properly?

- If **yes**, go to the step.
- If **no**, restore the ECS server by shutting the server down and bringing up the clone.

F

ECS 3.0 Shutdown and Startup Procedures

This appendix provides the procedures to properly shutdown and startup the ECS 3.0 system.

In This Appendix

Shutdown the ECS 3.0 System

This section provides the procedures to shutdown VMs within the ECS 3.0 system and includes the order in which you must shutdown the system. Use one of the following methods to shutdown the ECS 3.0 nodes.

- *Shutting Down VMs Using vCenter* (on page 210)
- *Shutting Down VMs from the Command Line* (on page 211)

Shutting Down VMs Using vCenter

Important: Please pay careful attention to the steps as virtual machines must be shutdown and in a specific order.

Complete the following procedure to shutdown ECS 3.0 VMs if you have access to vCenter but do not have access to the VMs from a Console or a terminal window.

1. Access your vCenter application.
2. Right-click the appropriate **ECS** virtual machine and select **Shut Down Guest OS**.
3. When prompted to shutdown the guest operating system, click **Yes**.

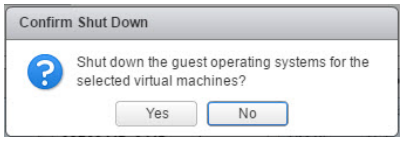

4. Monitor the **Recent Tasks** area to verify that the VM successfully shuts down.

   Important: Do not shutdown the next VM until the current VM completely shuts down and powers off.
5. Repeat Steps 1 through 4 to shutdown the remaining VMs, in the order listed below.
   - ECS node, if you have redundant ECS nodes
   - BOA/BST nodes
   - VCS Console nodes
   - Consul nodes
   - RAC nodes
   - Admin node, if necessary

Shutting Down VMs from the Command Line

Important: Please pay careful attention to the steps as virtual machines must be shutdown and in a specific order.

Complete the following procedure to shutdown ECS 3.0 VMs if you have access to the command line via a Console or a terminal window but do not have access to vCenter.

1. From a Console or terminal window, log into the ECS VM as **admin** user.
2. Enter the following command and press **Enter** to shutdown the VM.

   ```
   [admin@[VM] ~]$ sudo init 0
   ```

   Important: Wait until the VM completely shuts down and powers off before you shutdown the next VM in the system.
3. Repeat Steps 1 through 2 to shutdown the remaining VMs in the order listed below.
   - ECS node, if you have redundant ECS nodes
   - BOA/BST nodes
   - VCS Console nodes
   - Consul nodes
   - RAC nodes
   - Admin node, if necessary

Starting the ECS 3.0 System

Complete the following procedure, in order, to start the VMs in your ECS 3.0 system.

Notes:

- This section assumes that all ECS 3.0 nodes have been shutdown.
- The nodes can only be started from vCenter.
- If the Admin Node was also shutdown, you can restart this node at any time as it is independent of the ECS 3.0 nodes.

1 From vCenter, right-click the each Consul node, one by one, and select **Power On**.

2 After the Consul nodes boot, log into one of the Consul nodes as **admin** user.

3 Enter the following command to verify that the Consul nodes are up and running.

```
[admin@consul ~]$ consul members
```

```
[admin@consul ~]$ consul members
Node            Address              Status  Type    Build  Protocol  DC
.
.
consul1-test    172.20.167.94:8301   alive   server  0.7.2  2         dc1
consul2-test    172.20.167.95:8301   alive   server  0.7.2  2         dc1
consul3-test    172.20.167.96:8301   alive   server  0.7.2  2         dc1
.
.
```

4 Is the status for all three Consul nodes "**alive**"?

- If **yes**, go to the next step.
- If **no**, contact Cisco Services.

5 From vCenter, right-click the RAC nodes and select **Power On**.

6 After the RAC nodes boot, log into one of the nodes as **root** user.

7 Enter the following command to switch to **oracle** user.

```
# su - oracle
```

8 Enter the following command to set the CRS environment.

```
$ CRS
```

```
ORACLE_UNQNAME=CABHE
NLS_LANG=AMERICAN_AMERICA.AL32UTF8
LD_LIBRARY_PATH=/opt/oracle/installed/oracle_ee-12.1.0.2-0/lib32:/opt/oracle/installed/oracle_ee-12.1.0.2-0/lib
ORACLE_SID=CABHE01
ORACLE_BASE=/opt/oracle/orabase
ORAENV_ASK=NO
QTLIB=/usr/lib64/qt-3.3/lib
SHLIB_PATH=/opt/oracle/installed/oracle_ee-12.1.0.2-0/lib32:/opt/oracle/installed/oracle_ee-12.1.0.2-0/lib
ORA_NLS33=/opt/oracle/installed/oracle_ee-12.1.0.2-0/ocommon/nls/admin/data
ORACLE_HOME=/opt/oracle/installed/oracle_cluster-12.1.0.2-0
```

9 Enter the following command to execute a resource check.

```
$ crsctl status resource -t
```

```
--------------------------------------------------------------------------------
Name           Target  State        Server                   State details
--------------------------------------------------------------------------------
Local Resources
--------------------------------------------------------------------------------
ora.CRS.dg
               ONLINE  ONLINE       ecs-30-raca              STABLE
               ONLINE  ONLINE       ecs-30-racb              STABLE
ora.LISTENER.lsnr
               ONLINE  ONLINE       ecs-30-raca              STABLE
               ONLINE  ONLINE       ecs-30-racb              STABLE
ora.ORABACK.ORABACK_VOL.advm
               ONLINE  ONLINE       ecs-30-raca              Volume device /dev/a
                                                             sm/oraback_vol-73 is
                                                             online,STABLE
               ONLINE  ONLINE       ecs-30-racb              Volume device /dev/a
                                                             sm/oraback_vol-73 is
                                                             online,STABLE
.
.
--------------------------------------------------------------------------------
Cluster Resources
--------------------------------------------------------------------------------
ora.LISTENER_SCAN1.lsnr
      1        ONLINE  ONLINE       ecs-30-racb              STABLE
ora.LISTENER_SCAN2.lsnr
      1        ONLINE  ONLINE       ecs-30-raca              STABLE
ora.MGMTLSNR
      1        ONLINE  ONLINE       ecs-30-raca              169.254.237.195 192.
                                                             168.1.2,STABLE
ora.cabhe.cabhe_connection.svc
      1        ONLINE  ONLINE       ecs-30-raca              STABLE
.
.
```

10 Are all of the services online and stable?

- If **yes**, go to the next step.
- If **no**, contact Cisco Services.

11 From vCenter, select the VCS Console nodes, one by one, and select **Power On**.

12 In a Web browser, enter the following command to verify that the VCS Console is running.

URL Format:

```
https://[VCS Console IP]:6605/vcsconsole
```

Example:

```
https//10.90.47.21:6605/vcsconsole
```

13 Click the **Address** column header and sort the IP addresses in ascending or descending order.

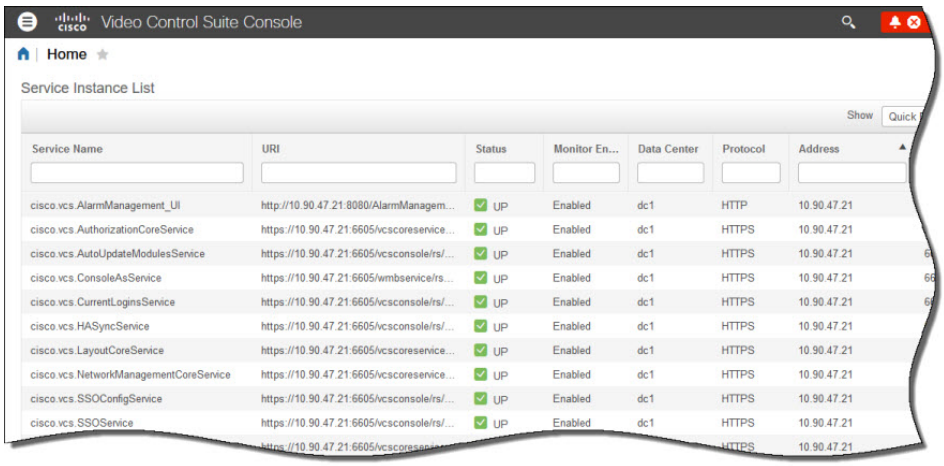

14 Scroll to the IP address for your primary VCS Console and verify that all services are **UP**.

15 Scroll to the IP address for your secondary VCS Console and verify that all services are **UP**.

16 Is the status for each VCS Console service **UP**?

- If **yes**, go to the next step.
- If **no**, contact Cisco Services.

17 From vCenter, select the ECS and the BOA/BST nodes, one by one, and select **Power On**.

Note: You can start up the ECS and BOA/BST nodes in any order.

18 From the VCS Console, verify that the status for all ECS and BOA/BST services are **UP**.

19 Is the status for each ECS and BOA/BST node **UP**?

- If **yes**, go to the next step.
- If **no**, contact Cisco Services.

20 Go to ***Verifying ECS Functionality*** (on page 140).

G

Consul Server Recovery

ECS 3.0 includes 3 Consul servers distributed in 2 UCS-M chassis. This means that two of the three Consul servers will be located in one of the two chassis.

If for some reason the chassis that contains the two Consul servers goes down, the entire communication between ECS services no longer be available. This is because the Consul cannot operate with a single Consul server (i.e. there is no quorum).

This appendix describes how you can recover a Consul server that is built on a failed ESXi host.

Important: Only follow this procedure if chassis A cannot be recovered. This procedure also be applies to any of the VMs under the chassis affected.

In This Appendix

Scenario Depicting a Failed Chassis With Two Consul Nodes

This section includes an example of a failed chassis that hosts two of the three Consul nodes in your NextX system. In this scenario, Chassis A goes down. As a consequence, the whole ECS system is down because the services cannot communicate to each other with only one Consul node (i.e. Consul3) running.

Note: Although all the VMs hosted on Chassis A are down, the VM storage still exists in the Guest VM storage. If the VM storage is not corrupt, the VMs can be recovered on one of the ESXi hosts on Chassis B.

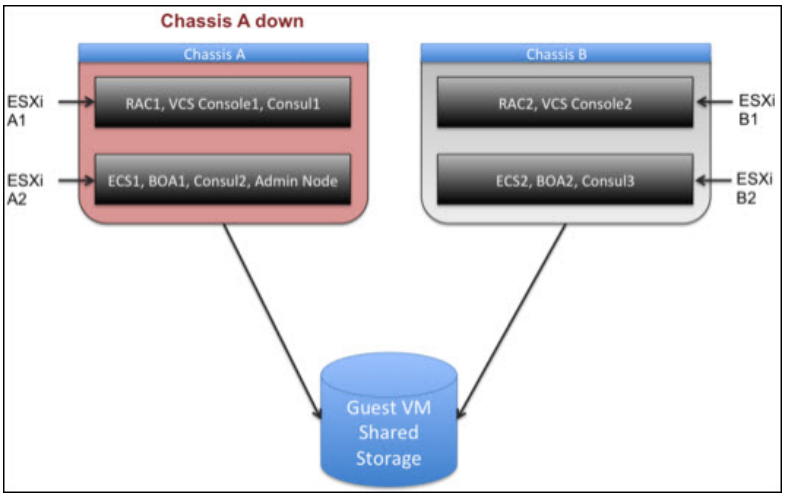

Recovering a Consul Server From a Failed ESXi Host

When Chassis A is down, all of the ESXi hosts and VMs still exist on the vCenter inventory. In order to re-add the Consul2 VM to the inventory on a different ESXi host (i.e. Chassis B), it must first to be removed from Chassis A.

Note: You could recover Consul1 VM; however, it does not make any difference which VM you recover from the failed host. The goal is to have at least two Consul servers running.

Complete the following procedure to recover a Consul Server from a failed ESXi host and move it to the working ESXi host.

1. From the vSphere Web UI, right-click the failed ESXi host and select **Disconnect**.
2. When prompted to disconnect the host, click **Yes**.
3. Monitor the **Recent Tasks** area to verify that the host was successfully disconnected.
4. Right-click the failed ESXi host again and select **All vCenter Actions > Remove** from Inventory. The Remove Host window appears.

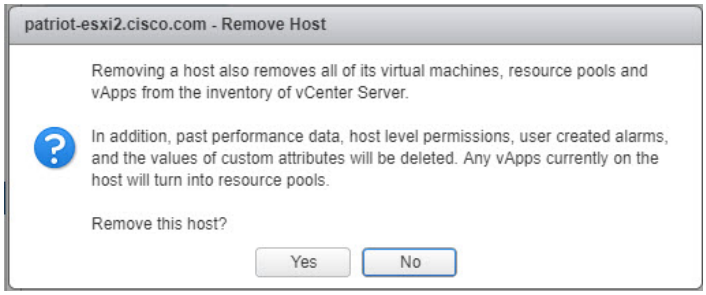

5. When prompted to remove the host, click **Yes**.

   Result: The ESXi host and its associated VMs are removed from the inventory of the vCenter Server.
6. Monitor the **Recent Tasks** area to confirm that the host was successfully removed from vCenter.

   Note: Once the Consul node that was on the failed VM is no longer in the vCenter inventory, it can be added to another ESXi host.
7. Complete the following steps to add one of the Consul nodes to another ESXi host.

   Note: This example uses Consul2 as the node that is moved to a new ESXi host.

   a. From the vSphere Web UI, click the **Home** icon, .

   b. From the Inventories area, click the **Storage** icon, Storage.

c From the left area of the window, locate and click the datastore where the Consul2 VM directory exists.

d From the right area of the window, click the **Manage** tab and then click the **Files** tab.

e Select the appropriate folder for the Consul2 files. The right area of the window displays the contents of the Consul2 folder.

Example:

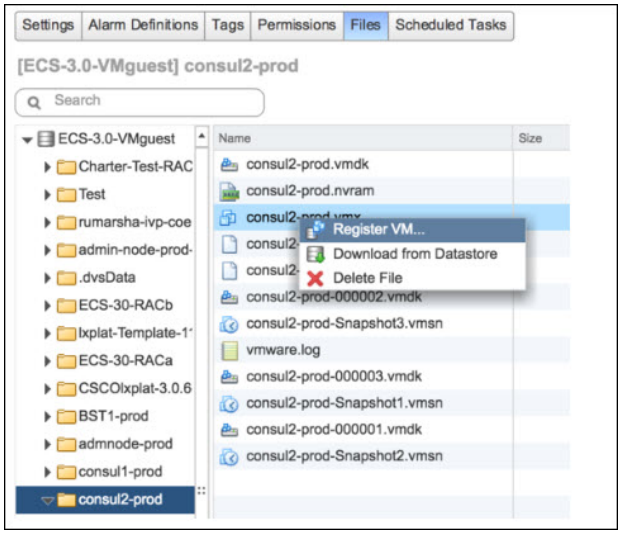

f Right-click the **vmx** file for Consul2 and select **Register VM**. The Register Virtual Machine > Name window displays.

g In the **Name** box, enter a name for the virtual machine.

h Select a folder where the Consul2 VM will be located.

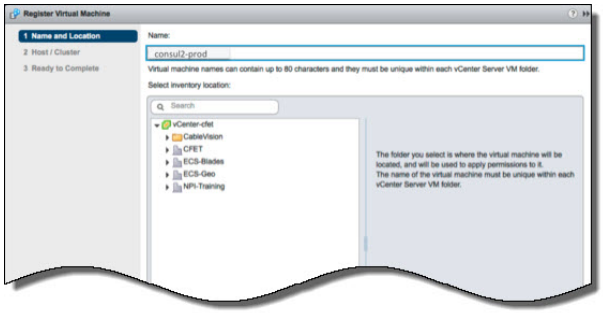

i Click **Next**. The Host/Cluster window displays.

j From the **Select inventory location** area, select an ESXi host from Chassis B where you want the Consul2 VM to reside. The Ready to Complete window displays.

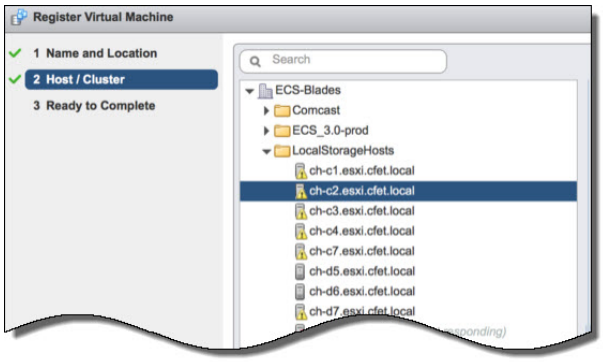

k Review the settings and then click **Finish**.

l Monitor the **Recent Tasks** area until the Consul2 VM is successfully registered to the new ESXi host.

8 Go to the ESXi host where the Consul2 VM resides and right-click the the **Consul2 VM**.

9 Select **Power On**. The following message will display.

10 Within the warning, click **Answer Question**. The following window displays.

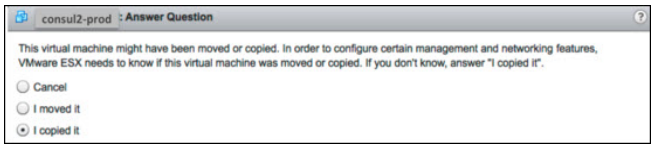

11 Click **I copied it** and then click **OK**.

12 Once the Consul2 node boots up, log into the VM as **admin** user.

13 Go to *Verifying Consul Functionality* (on page 53).

H

Procedures When Using an ESXi Client

This appendix describes various procedure when deploying and reconfiguring VMs using the a vSphere client.

In This Appendix

Deploy and Configure a VM Using an ESXi Client

Deploying a Virtual Machine Using vSphere Client

Complete the following steps to deploy a VM from an ESXi client for your NextX ECS 3.0 system.

1 Log on to the ESXi client.

2 From the File menu, select **Deploy OVF Template**. The Source window displays.

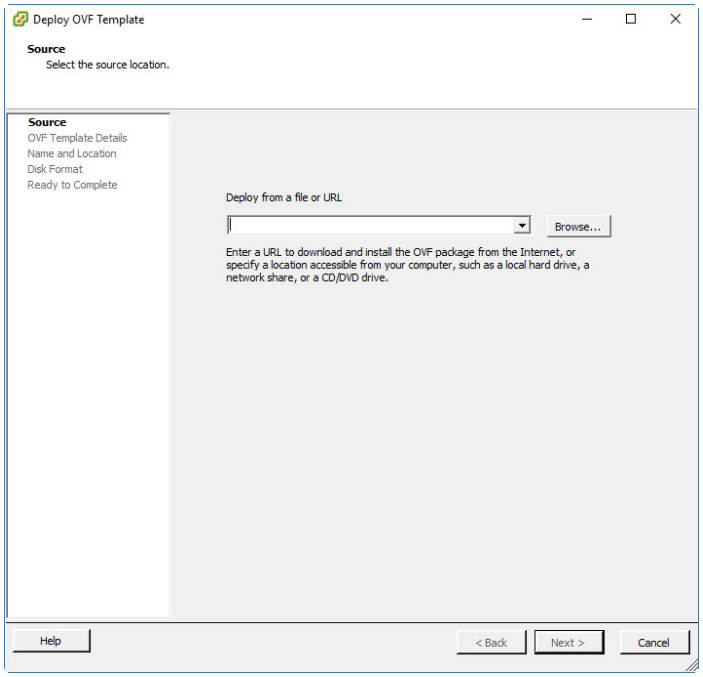

3 Click **Browse** and navigate to the directory where the Cisco platform OVA (i.e. CSCOlxplat-3.0.6.ova) resides.

4 Select the OVA and click **Open**. The absolute path to the OVA is added to the text box in the Source Window.

5 Click **Next**. The OVF Template Details window displays.

6 Review the details and click **Next**. The End User License Agreement displays.

7 Review the license agreement and click **Accept**. Then click **Next**. The Name and Location window display.

8 In the **Name** text box, enter a name to describe the VM. Then click **Next**. The Deployment Configuration window displays.

9 From the **Configuration** dropdown menu, select the **4CPU 4RAM 20GB** disk configuration and click **Next**.

10 From the **Select virtual disk format** drop-down menu, select **Thick Provision Lazy Zeroed** and click **Next**.

11 Select the appropriate **Destination Network** for each Source Network. Then click **Next**. The Ready to Complete window displays.

12 Review the configuration and then click **Finish**. A Deploying window appears and remains open until the new virtual machine is deployed.

Reconfiguring the Virtual Network Using vSphere Client

Complete the following steps to reconfigure the virtual hardware for the SR VM using vSphere client.

Use the following table to update the configuration for each respective VM in your system.

	Memory (GB)	CPUs	Hard Disk 1 (GB)
Consul	2	1	32
VCS Console	16	8	160
ECS	4	4	32
BOA	4	4	32

1 Select and right-click the new SR VM. The Edit Settings window displays.

2 Click **Memory** and modify the entry to the value defined for the VM in the above table.

Note: Make sure you change MB to **GB**.

3 Click **CPUs** and then from the **Number of virtual sockets** dropdown list, modify the entry to the value defined for the VM in the above table.

4 Select **Hard disk 1** and modify the **Provisioned Size** to modify the entry to the value defined for the VM in the above table.

5 Are you reconfiguring a Consul node?

- If **yes**, you have completed this procedure.
- If **no**, go to the next step.

6 Is the Oracle RAC set up on a different network?

- If **yes**, you have completed this procedure.
- If **no**, go to the next step.

7 Click the **Add** button and select **Ethernet Adapter**. Then click **Next**.

8 Maintain the VMXNET 3 adapter type and from the **Network** label dropdown menu, select the label associated with the RAC database. Then click **Next**. The Ready to Complete window displays.

9 Review the details and then click **OK**. The VM is reconfigured.

10 Review the options and then click **Finish**. The new Ethernet adapter is added to the virtual machine list.

Index

D

E

H

I

M

P

R

S

T

U

V

X

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA

http://www.cisco.com
Tel: 408 526-4000
800 553-6387
Fax: 408 527-0883

September 2017

Part Number TP-00133-01