ERICSSON

Small Integrated ENM System Administrator Guide

Operating Instructions

1/1543-CNA 403 3456 Uen BN

Contents

1 Introduction

This document describes the system administration tasks and troubleshooting guide for Small Integrated ENM deployments.

System Administration Tasks

The purpose of this document is to provide step-by-step instructions on how to configure and troubleshoot Small Integrated ENM deployments.

Target Audience

The intended audience of this document is system administrators who configure and manage Small Integrated ENM deployments.

Typographic Conventions

For more information on the typographic conventions used in this document, refer to the ENM Typographical Conventions.

Note: Be careful if you are cutting and pasting content from code examples. Sometimes the line spacing is altered and some examples might have actual values rather than variables.

2 Small Integrated ENM System Administration

This section describes the system administration tasks for a Small Integrated ENM deployment.

3 General Prerequisites

This section covers prerequisites that are essential for performing various administrative procedures described in this document.

3.1 Small Integrated ENM and VNF-LCM Security

Certain administrative procedures detailed below require access to either the internal, external, or both the internal and external VNF-LCM network.

— Access to the VNF-LCM internal network must be provided to software running on the Virtual Management Server(VMS).

Note: For instructions on how to enable VMS to access the VNF-LCM's internal network, see the section *Allow Access for ENM Whitelist VMs Over Internal Network to VNF-LCM* in *ENM on Cloud Deployment Instructions (2/153 72-AOM 901 151)*.

— Access to the VNF-LCM external network must be provided to web browsers running on various client machines.

Note: For instructions on how to manage access to VNF-LCM's external network, see the subsection *VNF-LCM Admin CLI* (specifically, sub subsections dealing with the *VNF-LCM Security Utility*) in ENM Configuration System Administrator Guide.

3.2 Browser Requirement

Some of the procedures in the later sections require access to vSphere management interfaces. For more information on client requirements, refer to the *Client Software Requirements* section of the *vSphere product documentation*.

4 Platform Security Hardening

This section provides procedures to enhance security features for the Small Integrated ENM platform.

Prerequisites

— vCenter is installed.

— Access to the vCenter client and the administrator user credentials are known.

— *ENM on Cloud Site Engineering Data (2/1057-AOM 901 151)* is available.

4.1 Disable SSH Service in vCenter

This section provides the procedure to disable SSH service to vCenter.

Steps

1. Log on to the vCenter Server Appliance Management Interface at `<vcenter_ip_vio_mgt>:5480` as user `administrator@vsphere.local` with password `<vcenter_sso_password>` from the ENM SED.
2. On the left plane, select **Access**.
3. On the right plane, click **EDIT**.
4. Uncheck the **Enable SSH login** check box.
5. Click **OK**.

Results

SSH service to vCenter is disabled.

4.2 Enable SSH Service in vCenter

This section provides the procedure to enable SSH service to vCenter.

Note: This procedure is only performed when facilitating the automated install and upgrade of the Small Integrated ENM platform.

Steps

1. Log on to the vCenter Server Appliance Management Interface at `<vcenter_ip_vio_mgt>:5480` as user `administrator@vsphere.local` with password `<vcenter_sso_password>` from the ENM SED.
2. On the left plane, select **Access**.
3. On the right plane, click **EDIT**.
4. Check the **Enable SSH login** check box.
5. Click **OK**.

Results

SSH service to vCenter Server is enabled.

4.3 Disable SSH Service in ESXi Hosts

This section provides the procedure to disable SSH service to the ESXi hosts.

Steps

1. Log on to the vCenter Client at `<vcenter_ip_vio_mgt>/ui/` as user `administrator@vsphere.local` with password `<vcenter_sso_password>` from the ENM SED.
2. Navigate to **Hosts and Clusters** (keyboard shortcut CTRL + ALT + 2).
3. Select the cluster inventory object `<vcenter_object_prefix>_CLUS` in the **Navigator** pane and open the **Hosts** tab.
4. Click each ESXi host in the list and complete the following substeps.
    a. Click **Configure**.
    b. Select **System**.

c. Select **Services**.

d. Click **SSH**.

e. Click **Edit Startup Policy**.

f. Choose the **Start and stop manually** option from the **Startup Policy list**.

g. Click the **Stop** button.

h. Click **OK**.

Results

SSH service to all ESXi hosts is disabled.

4.4 Enable SSH Service in ESXi Hosts

This section provides the procedure to enable SSH service to the ESXi hosts.

Note: This procedure is only performed when facilitating the automated install and upgrade of the Small Integrated ENM platform.

Steps

1. Log on to the vCenter client at `<vcenter_ip_vio_mgt>`/ui/ as user `administrator@vsphere.local` with password `<vcenter_sso_password>` from the ENM SED.
2. Navigate to **Hosts and Clusters** (keyboard shortcut CTRL + ALT + 2).
3. Select the cluster inventory object `<vcenter_object_prefix>_CLUS` in the **Navigator** pane and open the **Hosts** tab.
4. Click each ESXi host in the list and complete the following substeps.
    a. Click **Configure**.
    b. Select **System**.
    c. Select down to the **Services** section.
    d. Click **SSH**.
    e. Click **Edit Startup Policy**.
    f. Choose the **Start and stop manually** option from the **Startup Policy list**.
    g. Click the **Start** button.

Results

SSH service to all ESXi hosts is enabled.

5 Monitoring the VIO Stack

The section describes the tasks required to access the monitoring facility for VMware Integrated OpenStack (VIO).

Prerequisites

— ESXi has been successfully deployed and vSphere GUI is accessible.

— The VIO vApp has been deployed and configured.

— The `System Monitoring` user has been created in vCenter.

Steps

1. Log on to the vCenter client at `<vcenter_ip_vio_mgt>`/ui/ as user **`administrator@vsphere.local`** with password **`<vcenter_sso_password>`** from the ENM SED.
2. Click the **vCenter Server** instance in the **Navigator** pane to expand it and continue expanding objects until the **VIO-Openstack Resource Pool** and **VIO vApp** can be seen.

Results

The **Monitoring** console is accessible to the user.

5.1 Sample Use Cases

— To monitor the VIO-OpenStack Resource Pool, click the **Resource Pool** object, then in the center pane, select **Monitoring**.

  Refer to *VMware vCenter Monitoring* documentation for details on usage.

— To monitor the VIO vApp, click the **VIO vApp** object, then in the center pane, select **Monitoring**.

  Refer to *VMware VIO Monitoring* documentation for details on usage.

— To monitor Cluster Resource Usage, click the **Cluster** object, then in the center pane, select **Monitoring**. Click the **Performance** button, then ensure **Overview** is selected in the menu to the left.

  The graphs showing usage may need to be clicked to activate. `CPU` and `RAM` usage for the cluster can be seen in the graphs.

Additional Information

Further information on VIO monitoring can be found here: https://blogs.vmware.com/openstack/vmware-integrated-openstack-monitoring/

Further information on vCenter monitoring can be found here: https://docs.vmware.com/en/VMware-vSphere/6.7/vsphere-esxi-vcenter-server-67-monitoring-performance-guide.pdf

6 Restore OpenStack Services

This section provides steps to recover the OpenStack services in there is a failure.

Prerequisites

— VIO is installed.

— Small ENM SED in JSON format is populated and available on VMS.

Overview

— Recover OMS connectivity to VIO on page 10.

— Restart OpenStack Deployment on page 11.

— Shut down and Restart the VIO OMS vApp from the vCenter Web Client on page 12.

Result

All the OpenStack services are recovered and VIO is in healthy state.

6.1 Recover OMS connectivity to VIO

This section provides steps to re-connect OMS to VIO in vCenter.

Prerequisites

— VIO is installed.

— Management server `OMS` is up and running.

— Small ENM SED, in JSON format, is populated and available on VMS.

Steps

1. Log on to the vCenter Client at https://`<vcenter_ip_vio_mgt>`/ui/ as user `administrator@vsphere.local` with password `<vcenter_sso_password>` from the ENM SED.

2. Click on the **Menu** icon at the top of the vCenter Client and select **VMWare Integrated OpenStack**.
3. Go to **Basic Tasks** > **Connect to an OpenStack management server**.
4. Select **management-server** > **OK**.
5. On **View Certificate**, click **OK**.
6. Monitor the connecting progress, it takes approximately 5-10 minutes.

Results

OMS is connected to VIO in vCenter.

6.2 Restart OpenStack Deployment

> This section provides steps to restart the OpenStack deployment in vCenter.

Prerequisites

— VIO is installed.

— Small ENM SED, in JSON format, is populated and available on VMS.

Steps

1. Log on to the VMS as `root` user.
2. Log on to OMS server as `root` user.

```
[root@vms ~]# ssh viouser@oms
viouser@oms:~# sudo -i
```

3. Start all OpenStack services.

```
root@oms:~# viocli services start
```

Note: OpenStack deployment takes approximately 5-10 minutes to start.

Results

OpenStack deployment is started in vCenter and VIO is in healthy state.

6.3 Shut down and Restart the VIO OMS vApp from the vCenter Web Client

> This section provides the steps to shut down and restart the VIO OMS vApp from the vCenter client.

Prerequisites

— VIO is installed.

— `Connection Failed` error when checking the status of VIO in the vCenter client.

Steps

1. Shut down the VIO OMS vApp from the vCenter client.
    a. Log on to the vCenter client at https://`<vcenter_ip_vio_mgt>`/ui as user `administrator@vsphere.local` with password `<vcenter_sso_password>` from the SED.
    b. Select **Menu > Hosts and Clusters** in the vCenter Client.
    c. Right-click the vApp containing VM `management-server` in **Navigator** and select **Power > Shut Down**.
    d. Wait for the OMS vApp to shut down.
2. Start up the VIO OMS vApp from the vCenter client.
    a. Select **Menu > Hosts and Clusters** in the vCenter Client.
    b. Right-click the vApp containing VM `management-server` in **Navigator** and select **Power > Power On**.
    c. Open a VM console and wait for the `management-server` VM to start up.

Results

OMS is connected to the VIO in vCenter.

7 Small Integrated ENM Health Check

This section provides the steps to perform a Small ENM health check.

7.1 Perform Platform Health Check

This section provides steps to perform a platform health check.

Note: Health check is not supported in maintenance mode.

Prerequisites

All platform components installed.

Steps

1. A known issue triggered by a missed cleanup following a cancelled Backup or Snapshot Deployment workflow leaves snapshot disks in the system that are not visible in vCenter. The issue causes VNF-LCM workflows that detach volumes such as ENM High Availability, ENM Upgrade, and Manager ENM - shut down to fail. To avoid the issue always make sure the documented cleanup procedure is run immediately after a canceled Backup or Snapshot Deployment workflow. The cleanup must be run before either of the workflow runs again.

   Search for the issue as follows:

   a. Log on to the VMS as user **root**.

   b. From the VMS, SSH to the ESXi host with IP **<esxi_host1_ip_vio_mgt>** as user root with password **<esxi_host1_mgt_password>**.

   c. Change directory to the volume datastore.

      — For a SIENM Multi-Technology deployment:

      ```
      [root@esxi_1]# cd /vmfs/volumes/vsanDatastore
      ```

      — For a SIENM Transport Only deployment:

      ```
      [root@esxi_1]# cd /vmfs/volumes/datastore1
      ```

   d. Search for flag `ddb.deletable = "false"` in volume VMDK files as described below:

```
[root@esxi_1]# grep -ir "ddb.deletable.*false" */*[\)].vmdk 2>/dev/ →
null | grep -iv "image" | grep -iv ERICvms
[root@esxi_1]#
```

Contact Ericsson support for the issue recovery if search results are found.

Note:
- Major workflows such as ENM Upgrade and Manage ENM must not run until the issue recovery is complete.
- The recovery steps do not require an ENM application outage as long as no major workflow is running.

2. Check vSAN cluster health.

   Note: Perform this step if the deployment is Small Integrated ENM Multi-Technology.

   a. Go to **Navigator** and select the cluster inventory object **<vcenter_object_prefix>_CLUS**.

   b. Select the **Monitor** tab and click **vSAN**.

   c. Select **Skyline Health**.

   d. Expand each vSAN test and resolve the failed test issues apart from the following:

      - **Hardware compatibility** tests
      - **Disk format version** (This test can be disabled by clicking **Disk format version** > **Silence Alert** > **Yes**)
      - **vSAN Build Recommendation** (This test can be disabled by clicking **vSAN Build Recommendation** > **Silence Alert** > **Yes**)
      - **vSAN Disk Balance**

Note: If the health check **vSAN object health** has failed, click the **Repair Objects immediately** button, and then click the **Retest** button to refresh the test result.

- Select **vSAN** > **Resyncing Components** and confirm that there are no resyncing components.
- Select **vSAN** > **Virtual Objects** and confirm vSAN Object Health is healthy for all VMs.

  Ignore if the vSAN Object Health of `<vcenter_object_prefix>_VCSA VM` **VM Home** is `Non-availability` related.
- Select **vSAN** > **Physical Disks** and confirm that all disks are healthy.

3. Confirm all VMs are compliant with vSAN storage policy FTT=1.

   Note: Perform this step if the deployment is Small Integrated ENM Multi-Technology.

   a. Go to **Navigator** and select the cluster inventory object **<vcenter_object_prefix>_CLUS**.

   b. Select the **VMs** tab.

   c. Confirm that column **VM Storage Policies Compliance** is compliant for all VMs.

      Note:
      - If the column is not displayed, left-click the down arrow on one of the column bar and select **Show/Hide Columns** > **VM Storage Policies Compliance**.
      - Click the column to sort by **VM Storage Policies Compliance** and again to invert the sort (all VMs are compliant if the first row of both sorts shows compliant).
      - Ignore if `<vcenter_object_prefix>_VCSA` is the only VM in `Non-compliant` state for **VM Storage Policies Compliance** and the vSAN Object Health of `<vcenter_object_prefix>_VCSA VM` **VM Home** is `Non-availability` related.

   d. Multi-select any VMs that show **Out of Date** compliance if necessary, right-click the VM selection and select **VM Policies** > **Reapply VM Storage Policy**.

Note:
- Any VM with text (orphaned) next to its inventory object in the **Navigator** does not block further activities.
- To remove orphaned VMs, follow the procedure Remove Orphaned VMs from vCenter Inventory on page 101.

4. Check the remaining components by logging on to VMS at `<vms_ip_vio_mgt>` as `root` user with password `<vms_root_password>` and running the following command.

```
[root@vms]# /opt/ericsson/senm/bin/sienm_hc.sh
```

Example

```
[root@vms]# /opt/ericsson/senm/bin/sienm_hc.sh


--------------------------------------------------------------------------------
-----------------

Small Integrated ENM Platform Health Check

--------------------------------------------------------------------------------
-----------------


15:08:12 INFO Deployment type is: Multi-Technology

15:08:12 INFO Checking ESXi
...

15:10:21 INFO Completed with errors and warnings

15:10:21 INFO The Summary file generated at: /vol1/senm/log/html/HC_Report_2
020_01_16_15_08_12.yml

15:10:21 INFO Log file used: /vol1/senm/log/hc_log/HC_Log_2020_01_16_15_08_1
2

15:10:21 INFO The HTML Report file generated at: /vol1/senm/log/html/HC_Repo
rt_2020_01_16_15_08_12.html
```

Note: Created reports are available in the `/vol1/senm/log/html/` directory.

5. View the report by transferring the html file to a machine running a web browser.

Note: The report can also be viewed directly on the VMS using `elinks` as shown below.

```
[root@vms~]# elinks <report_path>
```

Example

```
[root@vms~]# elinks /vol1/senm/log/html/HC_Report_2020_01_16_15_08_12.html
```

6. Follow the instructions within the report.

Stop!

Any errors reported must be resolved before proceeding with the planned activity. If only warnings are reported, continue with the planned activity and investigate at a later date.

7.2 Platform Health Check Remedial Actions

> This section provides remedial actions to any issues detected by the health check.

If the health check reports anything that is not covered in this section, contact Ericsson Support.

7.2.1 ESXi Hosts

Connection State

This could be caused by loss of power or a hardware failure. Investigate probable causes and refer to section Unplanned Loss of ESXI Host on page 44.

Contact Ericsson Support in the case of any difficulties.

ESXi Version / NICs Driver / NICs Firmware / RAID Controller Version

Verify that all firmware and software versions are the same across all hosts. Refer to the document *FLARE and Firmware Handling Guide for HP/EM*.

Inode Usage / Space Usage

1. Log on to the ESXI host using `<esxi_host<host_number>_ip_vio_mgt>` as user `root` with password `<esxi_host<host_number>_mgt_password>`.
2. Check the amount of free space on the file system.

   ```
   [root@esxi:~] localcli system visorfs ramdisk list
   ```

3. Ensure that the root file system has at least 20% free space.
4. Ensure that there are free inodes on each file system by comparing the `Maximum Inodes` with the `Used Inodes`.

Note: If there are no free inodes on the var file system, clean up the files in directory `/var/run/vmware/tickets`.

```
[root@esxi:-] cd /var/run/vmware/tickets/
[root@esxi:-] find . -name 'vmtck-*' | xargs rm
```

NICs Status

This could indicate a hardware or cabling failure. Investigate probable causes and contact customer support.

RAID Controller Status

Health Check may report RAID Controller version as `NOK` occasionally if the ESXi Host does not receive the necessary information from the HW.

The RAID Controller version is the Smart Array Firmware Version and can be verified as described in section *Smart Array Firmware* in FLARE and Firmware Handling guide for HP/EMC.

Note: If the RAID controller firmware version is as expected, then the Health Check status for Gen9 RAID Controller version as `NOK` can be safely ignored.

7.2.2 Operations Management Server

Used space on /

1. Log on to the VMS as `root` user.
2. Log on to OMS server as `root` user.

```
[root@vms ~]#  ssh viouser@oms
viouser@oms:~#  sudo -i
```

3. Clean up the ansible files.

```
root@oms:~# find /tmp/.ansible/* -mtime +3   -exec rm -rf {} \;
```

Hanging Mount

1. Log on to the VMS as `root` user.
2. Log on to OMS server as `root` user.

```
[root@vms ~]#  ssh viouser@oms
viouser@oms:~#  sudo -i
```

3. Unmount the mount point.

```
root@oms:~# umount -l /mnt/backup
```

4. Remove the following entry from file `/etc/fstab` if it exists.

```
<vms_ip_vio_api>:/vol1/oms /mnt/backup nfs proto=tcp,port=2049 0 0
```

Critical Services

1. Log on to the VMS as `root` user.
2. Log on to OMS server as `root` user.

```
[root@vms ~]#  ssh viouser@oms
viouser@oms:~#  sudo -i
```

3. Restart any failed service.

```
root@oms:~#  systemctl restart <failed services>
```

Note: Where `<failed services>` are failed services listed in the health check report.

Deployment Status

If any of the OpenStack services are in failed state, restart all services.

1. Log on to the VMS as `root` user.
2. Log on to OMS server as `root` user.

```
[root@vms ~]#  ssh viouser@oms
viouser@oms:~#  sudo -i
```

3. Restart all OpenStack services.

```
root@oms:~#  viocli services stop
root@oms:~#  viocli services start
```

7.2.3 Load Balancer

Hanging Mount

1. Log on to the VMS as `root` user.

2. Log on to `loadbalancer01` as `root` user and unmount the mount point.

```
[root@vms~]# ssh viouser@oms
viouser@oms:~# ssh loadbalancer01
viouser@loadbalancer01:~# sudo -i
root@loadbalancer01:~# umount -l /mnt/backup
```

3. Remove the following entry from file `/etc/fstab` if it exists.

```
<vms_ip_vio_api>:/vol1/oms /mnt/backup nfs proto=tcp,port=2049 0 0
```

Services

Analyze the failure cause and restart the failed service.

Note: For help identifying the cause of the failure contact Ericsson Support.

1. Log on to the VMS as `root` user.
2. Log on to `loadbalancer01` as `root` user and restart the failed services.

```
[root@vms ~]# ssh viouser@oms
viouser@oms:~# ssh loadbalancer01
viouser@loadbalancer01:~# sudo -i
root@loadbalancer01:~# systemctl restart <failed services>
```

Note: Where `<failed services>` are services listed in the health check report.

HAProxy Status

Examine HAProxy logs on the load balancer and fix any issues.

```
[root@vms ~]# ssh viouser@oms
viouser@oms:~# ssh loadbalancer01
viouser@loadbalancer01:~# sudo -i
root@loadbalancer01:~# less /var/logs/haproxy/haproxy.log
```

Note: For help identifying the cause of the failure contact Ericsson Support.

7.2.4 OpenStack Compute

nova-compute service

Restart `nova-compute` service on `compute01` node to bring the system back up and contact Ericsson Support.

1. Log on to the VMS as `root` user.

2. Log in to `compute01` node as `root` user and restart `nova-compute`.

```
[root@vms ~]# ssh viouser@oms
viouser@oms:~$ ssh viouser@compute01
viouser@compute01:~$ sudo -i
root@compute01:~#systemctl restart nova-compute
root@compute01:~#
```

7.2.5 Virtual Management Server

Used space on /

Analyze space usage and clean up any unnecessary files.

Used space on /vol1

Clean up old artifacts on VMS (refer to section VMS Artifact Cleanup Procedure on page 102

7.2.6 VCenter Alarms

vCenter Alarms raised

1. Log on to the vCenter client URL https://<vcenter_ip_vio_mgt>/ui/ using `administrator@vsphere.local` with password `<vcenter_sso_password>`.
2. Open the **Alarms** tab.
3. Investigate any critical alarms.

7.2.7 Orphaned Resources

Note: If the health check reports some parts of this section as skipped, ensure that the `openstack .rc` files exist in `/vol1/senm/etc` on VMS.

The results may be invalid if any VNF-LCM workflow is in progress. Rerun the health check after the workflow has completed.

Instances / Virtual Machines / Shadow VMs

Refer to section Remove Orphaned VMs from vCenter Inventory on page 101.

Snapshots

Refer to section Delete Orphaned VM Snapshots on VIO (Method 2) on page 123.

Ports

1. Log in to VMS as a `root` user.
2. Delete any orphaned ports found.

```
[root@vms~]# openstack port delete <port ID(s)>
```

Note: Where `<port ID(s)>` are the IDs listed in the health check report.

7.3 ENM Health Check

> This section provides steps to perform an ENM health check.

Note: This section is not valid at initial installation.

Prerequisites

ENM is installed.

Steps

1. Check for running ENM HA workflows.
    a. Access the workflows using the following URL in a browser: http://`<external_ipv4_for_services_vm>`/index.html#workflows

    Note: Replace the value for `<external_ipv4_for_services_vm>` in the URL with the key defined in the VNF-LCM SED.

    b. Ensure that there are no HA workflows running.
2. Check consul members for any instances that not `Active`.
    a. Log on to the VMS on the management network `<vms_ip_vio_mgt>` as user `root` with password `<vms_root_password>`.
    b. Get the external IP address of the `emp` VM.

```
# openstack server list |grep emp
```

    c. Log on to `emp` VM and switch to root user.

```
# ssh -i /vol1/senm/etc/key-pair-vio-<vio_os_project_name>.pem clou
d-user@<emp external ip from above command>
[cloud-user@ieatvio5571-emp-0]$ sudo -i
[root@ieatvio5571-emp-0] #
```

d. Use the `consul members` command and check for any instance not `alive`.

Example

```
[root@ieatvio5571-emp-0 ~]# consul members | grep -v alive
Node                          Address           Status  Type
Build  Protocol  DC
```

Note: Ensure that there is no consul members in the list.

3. Check for any inconsistency of power state of VMs between vCenter and the consul.

a. Access the workflows using the following URL in a browser: http://<external_ipv4_for_services_vm>/index.html#workflows.

Note:
- Replace the value for <external_ipv4_for_services_vm> in the URL from the key defined in the VNF-LCM SED.
- If there is any instance of `Backup Deployment` or `Snapshot Deployment` running wait until it is finished.

b. Follow the section Delete Orphaned VM Snapshots on VIO (Method 2) on page 123 to delete any orphaned VM snapshots on vCenter.

c. Log on to the VMS on the management network <vms_ip_vio_mgt> as user `root` with password <vms_root_password>.

d. List any server instances in `SHUTOFF` state.

```
# openstack server list |grep SHUTOFF
```

e. If any instance is in `SHUTOFF` state and in `active` state in the `consul members list` from the previous step, then set the state to active.

```
# openstack server set --state active <server>
```

Note: Where <server> is the server (ID or Name) from the `openstack server list` command.

4. Log on to the VMS on the management network <vms_ip_vio_mgt> as user `root` with password <vms_root_password> and check that all VMs are in `ACTIVE` status using the following command.

Example

```
# openstack server list -c ID -c Name -c Status
+--------------------------------------+------------------------------+----
----+
| ID                                   | Name                         | Sta
tus |
+--------------------------------------+------------------------------+----
----+
| cd17956d-66e8-405f-86e6-7d6a40459b19 | dlv-c5a02-fmx-0              | ACT
IVE |
| 0b634e93-a514-44eb-900f-286fdf645c5f | dlv-c5a02-fmx-1              | ACT
IVE |
| f840f152-356e-4657-bf98-5a03ff7cac5c | dlv-c5a02-smrsserv-4         | ACT
IVE |
####### remaining output ommited #######
```

Note: Contact Ericsson customer support if all VMs do not have a status of `ACTIVE`.

8 Manual Recovery of an ENM Server Instance

This section provides steps on manually recovering an ENM server instance.

Manual recovery involves stopping ENM HA during the procedure.

Note:
— This procedure should only be run if ENM HA cannot automatically recover the instance.

— This procedure should only be run on direction from Ericsson support unless directly referenced from another documented procedure.

Steps

1. Disable ENM HA.

    a. Log on to the VMS on the management network `<vms_ip_vio_mgt>` as user `root` with password `<vms_root_password>`.

    b. Disable ENM HA.

```
[root@vms ~]# python /opt/ericsson/senm/lib/ha_tool.py set -s disabled
```

2. Open the VNF-LCM services GUI at URL http://`<external_ipv4_for_services_vm>`/index.html#workflows.

    Note: Where `<external_ipv4_for_services_vm>` is value from the VNF-LCM SED.

3. Browse to and cancel any HA workflow that is failing to recover the server instance to be manually recovered.

4. List the server instance to be manually recovered on the VMS as `root` user .

```
[root@vms ~]# openstack server list | grep <server instance>
```

5. Recover the server instance.

```
[root@vms ~]# /opt/ericsson/senm/utils/update_stack.sh -n <stack_name> -i nod →
e_index [-i node_index...]
```

Note:

— Where <stack_name> is the stack name in OpenStack without the prefix and suffix (for example, for stack ieatvio5567_ops_a3c1cce6-9a6a-4bb5-921c-4dbb10bb564b use ops).

— <node_index> is the server instance index (for example, for <deployment_id>-ops-1 use 1. The node_index option can be repeated to recover more than 1 instance.

Example

```
Example 1: To recover instance <deployment_id>-ops-0

[root@vms ~]# /opt/ericsson/senm/utils/update_stack.sh -n ops -i 0
14:24:32 INFO Get stack ID for stack ops...
14:24:32 INFO Get resource for stack ops...
14:24:33 INFO Find inner stack of NodeIdex: 0...
14:24:35 INFO Found ieatvio5567_ops_a3c1cce6-9a6a-4bb5-921c-4dbb10bb564b-ops →
-n7fmd3tunht6-0-nj4y7rq37o7a 14:24:35 INFO Get resource name...
14:24:36 INFO Found ops_definition_vm 14:24:36 INFO Mark resource unhealthy. →
..
14:24:36 DEBUG openstack stack resource mark unhealthy ieatvio5567_ops_a3c1c →
ce6-9a6a-4bb5-921c-4dbb10bb564b-ops-n7fmd3tunht6-0-nj4y7rq37o7a ops_definiti →
on_vm "Changed by update_stack"
14:24:37 DEBUG openstack stack resource list --filter type=OS::Nova::Server  →
ieatvio5567_ops_a3c1cce6-9a6a-4bb5-921c-4dbb10bb564b-ops-n7fmd3tunht6-0-nj4y →
7rq37o7a
+------------------+--------------------------------------+---------------- →
--+-----------------+----------------------+
| resource_name    | physical_resource_id                 | resource_type   →
  | resource_status | updated_time         |
+------------------+--------------------------------------+---------------- →
--+-----------------+----------------------+
| ops_definition_vm | e17deebe-c7ff-4536-908a-a61c253d992d | OS::Nova::Serve →
r | CHECK_FAILED | 2018-11-28T11:23:23Z    |
+------------------+--------------------------------------+---------------- →
--+-----------------+----------------------+
14:24:42 INFO Update inner stack...
14:24:42 DEBUG openstack stack update ieatvio5567_ops_a3c1cce6-9a6a-4bb5-921 →
c-4dbb10bb564b-ops-n7fmd3tunht6-0-nj4y7rq37o7a --existing --wait
2018-11-29 14:24:44Z [ieatvio5567_ops_a3c1cce6-9a6a-4bb5-921c-4dbb10bb564b-o →
ps-n7fmd3tunht6-0-nj4y7rq37o7a]: UPDATE_IN_PROGRESS Stack UPDATE started
2018-11-29 14:24:53Z [ieatvio5567_ops_a3c1cce6-9a6a-4bb5-921c-4dbb10bb564b-o →
ps-n7fmd3tunht6-0-nj4y7rq37o7a.ops_definition_vm]: CREATE_IN_PROGRESS state  →
changed
2018-11-29 14:25:19Z [ieatvio5567_ops_a3c1cce6-9a6a-4bb5-921c-4dbb10bb564b-o →
ps-n7fmd3tunht6-0-nj4y7rq37o7a.ops_definition_vm]: CREATE_COMPLETE state cha →
nged
2018-11-29 14:25:19Z [ieatvio5567_ops_a3c1cce6-9a6a-4bb5-921c-4dbb10bb564b-o →
ps-n7fmd3tunht6-0-nj4y7rq37o7a.ops_definition_volume_attach]: UPDATE_IN_PROG →
RESS state changed
2018-11-29 14:25:25Z [ieatvio5567_ops_a3c1cce6-9a6a-4bb5-921c-4dbb10bb564b-o →
ps-n7fmd3tunht6-0-nj4y7rq37o7a.ops_definition_volume_attach]: UPDATE_COMPLET →
E state changed
2018-11-29 14:25:32Z [ieatvio5567_ops_a3c1cce6-9a6a-4bb5-921c-4dbb10bb564b-o →
ps-n7fmd3tunht6-0-nj4y7rq37o7a]: UPDATE_COMPLETE Stack UPDATE completed succ →
essfully
+--------------------+------------------------------------------------------ →
--------------------------------+ |
Field                | Value                                                 →
                               |
+--------------------+------------------------------------------------------ →
--------------------------------+
| id                 | 4a5b3bab-1aa6-44f1-a0cc-441782bb1eb5 |
| stack_name         | ieatvio5567_ops_a3c1cce6-9a6a-4bb5-921c-4dbb10bb564b →
-ops-n7fmd3tunht6-0-nj4y7rq37o7a |
| description        | ops template                                          →
                               |
| creation_time      | 2018-11-28T11:23:23Z                                  →
                               |
| updated_time       | 2018-11-29T14:24:43Z                                  →
                               |
| stack_status       | UPDATE_COMPLETE                                       →
                               |
```

```
| stack_status_reason | Stack UPDATE completed successfully                    →
                                    |
+---------------------+------------------------------------------------------ →
--------------------------------+
```

Note: For `sentinel` and `servicereg` instances use the following `node_index` values.

Table 1 Sentinel and Servicereg node_index

Instance	node_index
sentinel	`0`
servicereg-1	`0`
servicereg-2	1
servicereg-3	2

Example

6. Monitor the server instance until the rebuilt server instance status is Example 2: To recover servicereg-3 server [root@vms ~]# /opt/ericsson/senm/utils/update_stack.sh -n serviceregistry -i 2`ACTIVE` and an internal IP address is assigned.

```
[root@vms ~]# openstack server list | grep <server instance>
```

7. Enable ENM HA.

```
[root@vms ~]# python /opt/ericsson/senm/lib/ha_tool.py set -s enabled
```

Results

The ENM server instance was successfully recovered.

9 Confirm VM Storage Policy Compliance with vSAN Storage Policy

This section details how to confirm that the VM storage complies with the vSAN storage policy.

1. Go to **Navigator** and select the cluster inventory object `<vcenter_object_prefix>_CLUS`.
2. Select the **VMs** tab.
3. Confirm that the column **VM Storage Policies Compliance** is **Compliant** for all VMs.

   Note:
   - If the column is not displayed, left-click the down arrow on one of the column bar and select **Show/Hide Columns** > **VM Storage Policies Compliance**.
   - Click the column to sort by **VM Storage Policies Compliance** and again to invert the sort. All VMs are compliant if the first row of both sorts shows **Compliant**.

4. Multi-select any VMs that show as **Non Compliant**, right-click the selection and select **VM Policies** > **Check VM Storage Policy Compliance**.
5. Multi-select VMs if they show **Out of Date** compliance, right-click the VM selection and select **VM Policies** > **Reapply VM Storage Policy**.

   Note:
   - Any VM with text `(orphaned)` next to its inventory object in **Navigator** does not block maintenance mode.
   - To remove orphaned VMs follow procedure Remove Orphaned VMs from vCenter Inventory on page 101

6. Monitor vSAN health checks in vCenter under **<vcenter_object_prefix>_CLUS > Monitor > vSAN > Skyline Health**.

Note:

— Exiting maintenance mode the vSAN may create mirror copies of large storage components on the returned host vSAN health check and **Data > vSAN object health** fails until the synchronization is complete as these components are not yet compliant with the vSAN storage policy.

— In vCenter, select **<vcenter_object_prefix>_CLUS > Monitor > vSAN > Virtual Objects** and sort by column **vSAN Object Health** to check for objects with `Reduced Availability`.

— Select and expand these objects to monitor the component syncing. The vSAN Data health checks pass when all the synchronization is complete.

10 Small Integrated Multi-Technology Host Startup, Shutdown, and Recovery

This section describes the maintenance and recovery procedures for Small Integrated ENM Multi-Technology deployments.

Note:
- Application restarts can be observed when a host is removed from the cluster for maintenance or upgrade or when a host is lost because of an unplanned outage. This can cause temporary loss of ENM application availability while the deployment is not running as a complete cluster. The length of the outage must not exceed the time it takes for ENM HA to restart the application.
- While the system is in maintenance mode, VM restarts can be observed.

10.1 Planned ESXi Host Maintenance

This section details how to complete a planned maintenance on an ESXi host.

Note:
- To allow continuous VM access, you can only place a single ESXi host in maintenance mode at a time.
- Before placing a host in maintenance mode, it is essential you first confirm the health of all vSAN storage objects on the remaining two hosts.

A Small Integrated ENM Multi-Technology deployment runs a three-host vSAN cluster for all VM storage with storage policy FTT=1 applied to all VM storage objects. All ESXi host maintenance procedures must handle the vSAN cluster policy correctly to maintain VM access during the maintenance interval.

vSAN storage policy FTT=1 ensures that a mirror copy of each VM storage object and a witness component is maintained on the other two hosts.

vSAN is an object-based distributed storage system that pools the directly attached disks from each ESXi host. VMs consist of various storage objects, for example, VM disks (VMDKs), VM home namespace, VM swap areas. Storage objects for a VM are distributed across all three hosts.

Since three hosts are required to apply the vSAN policy FTT=1 to each storage object (two copies plus one witness), virtual machines become non-compliant when a single host is placed in maintenance mode. Non-compliance means that there is no protection against the loss of a second host until the host exits

maintenance mode and its storage objects are synchronized from the other hosts. For this reason planned maintenance procedures must be as short as possible to reduce the risk of losing VM access if a second host is lost during the maintenance interval.

When the host exits maintenance mode, its storage objects are synchronized from the other two hosts until all objects are again compliant with FTT=1. Storage synchronization times increase with longer maintenance intervals.

10.1.1 ESXi Maintenance Mode with Healthy vSAN Cluster

> This section describes how to bring an ESXi host in and out of planned maintenance mode for a Small Integrated ENM Multi-Technology deployment.

Note:

- Replacement of faulty vSAN cache or capacity disks is not covered by this procedure.
- Only use this procedure when all VMs are compliant with the vSAN storage policy FTT=1. The steps to check vSAN storage object health are provided.
- You can only place a single host in maintenance mode at a time or VM access is lost.
- An ENM upgrade is not possible when a host is in maintenance mode.
- Maintenance mode is reached when all powered on VMs are migrated by vSphere to alternative hosts.

This procedure can be used for firmware upgrades, ESXi host patching, and most hardware replacements.

Do!

For a Small Integrated ENM Multi-Technology deployment:

— It is not possible to maintain ENM application availability if a host failure occurs when another host is already out of the cluster for maintenance. It is strongly recommended that you conclude any maintenance activities as quickly as possible and do not leave the deployment running on two hosts for more than a few hours.

— If all three hosts require a maintenance interval, leave the host running the `ControlStack` until last to avoid migrating the `ControlStack` to alternative hosts multiple times. The host running the `ControlStack` is the host currently running the vCenter appliance VM with name `<vcenter_object_prefix>_VCSA`, for example. `VIO_VCSA`. Select the VCSA VM in the vCenter **Navigator** pane. The host running the VM is displayed in the **Summary** tab in the middle pane.

Prerequisites

All VMs are compliant with the vSAN storage policy FTT=1.

Steps

1. Disable scheduled backups.

   Note: This step is not required if ENM has not been installed yet.

   — On deployments that use OMBS as the backup and restore storage solution: Follow the section *Deactivate Small Integrated ENM OMBS Policy* in Small Integrated ENM Backup and Restore System Administrator Guide

   — On deployments that use the customer-provided NFS Share as the backup and restore storage solution: Follow the section *Deactivate Small Integrated ENM Backup Policy to External Device or Customer-Provided NFS Share* in Small Integrated ENM Backup and Restore System Administrator Guide .

2. Log on to the vCenter client URL `https://<vcenter_ip_vio_mgt>/ui` as administrator `administrator@vsphere.local` with password `<vcenter_sso_password>`.

3. Check vSAN cluster health.

   a. Select the cluster inventory object `<vcenter_object_prefix>_CLUS` in **Navigator**.

   b. Select the **Monitor** tab and click **vSAN**.

   c. Select **Skyline Health**.

d. Expand each vSAN test and resolve the failed test issues apart from the following:

   — **Hardware compatibility** tests

   — **Disk format version** (This test can be disabled by clicking **Disk format version** > **Silence Alert** > **Yes**)

   — **vSAN Build Recommendation** (This test can be disabled by clicking **vSAN Build Recommendation** > **Silence Alert** > **Yes**)

   — **vSAN Disk Balance**

   Note: If health check **vSAN object health** has failed, click the **Repair Objects immediately** button and then click the **Retest** button to refresh the test result.

   — Select **Resyncing Components** and confirm that there are no resyncing components.

   — Select **Virtual Objects**. **vSAN Object Health** must be **Healthy** for all VMs

4. Follow the section Confirm VM Storage Policy Compliance with vSAN Storage Policy on page 28.

5. Confirm that there are no **ENM High Availability** workflows running.

   Note: This step is not required if ENM has not been installed yet.

   a. vSAN objectOpen a web browser to VNF-LCM services URL `http://<external_ipv4_for_services_vm>/index.html#workflows`.

      Note: Where `<external_ipv4_for_services_vm>` is the value from the VNF-LCM SED.

   b. Confirm row **High Availability Workflow**.

   c. Wait for any HA workflow instance to complete before entering maintenance mode.

6. Place the ESXi host in maintenance mode.

   Note: This step is not required if ENM has not been installed yet.

   a. Expand the cluster inventory object `<vcenter_object_prefix>_CLUS` in **Navigator** to display the ESXi hosts.

   b. Select the host to place in maintenance mode.

   c. Select the **VMs** tab in the middle pane to display all VMs associated with this host.

d. Right-click the host to be placed in maintenance mode and select **Maintenance Mode** > **Enter Maintenance Mode**.

e. Ensure that the **vSAN data migration** is set to **Ensure accessibility**.

f. Click **OK**.

g. Monitor the VM migrations that take place in **Recent Tasks** as vSphere looks for alternative hosts to maintain VM access.

h. Confirm the number of VMs on the host reduces as migrations take place in the **VMs** tab .

Note: This can take some time depending on load and if the host is running the `ControlStack` VMs.

i. Wait until only two VMs remain in the list.

Note: These are `<deployment_id>-neo4j-X` and `servicereg-X`, where `X` is the VM instance number running on the host entering maintenance mode.

7. Gracefully shut down `<deployment_id>-neo4j-X` and `servicereg-X`.

Note: This step is not required if ENM has not been installed yet.

a. Right-click `<deployment_id>-neo4j-X` in the **VMs** tab and select **Power** > **Shut Down Guest OS**.

b. Click **OK** on the confirmation dialog that appears.

c. Right-click `servicereg-X` in the **VMs** tab and select **Power** > **Shut Down Guest OS**.

d. Click **OK** on the confirmation dialog that appears.

Note: Take care not to select option **Power** > **Power Off** as this powers off a VM without a graceful shutdown of its operating system.

8. Wait a short interval for vCenter to display `(maintenance mode)` next to the host in the inventory.

Note: The host is now in maintenance mode.

9. Temporarily disable the DRS rule to achieve maintenance mode (optional: consider only if the host cannot enter maintenance mode)

Note: A DRS affinity rule exists to keep all ControlStack VMs running on the same host. Depending on load, it may become necessary to temporarily disable the DRS rule to achieve maintenance mode.

a. Navigate to **<vcenter_object_prefix>_CLUS>** > **Configure** > **Configuration** > **VM/Host Rules**.

b. Select rule `ControlStackAffinityRule`.

c. Click **Edit**.

d. Uncheck the **Enable rule** checkbox.

e. Click **OK**.

f. Wait a short interval for VM migrations to begin.

10. Disable vSphere DRS (optional: consider only if the host cannot enter maintenance mode).

    Note: vSphere DRS can be disabled to give the administrator full control over VM migrations which have to be done manually before the host can enter maintenance mode.

    a. Navigate to **<vcenter_object_prefix>_CLUS>** > **Configure** > **Services** > **vSphere DRS** > **Edit** > **Edit Cluster Settings**.

    b. For **Automation Level** select **Manual**.

    c. Click **OK** and follow the remaining substeps to manually migrate each powered-on VM to an alternative host with sufficient capacity .

    d. Right-click the VM in **Navigator** and select **Migrate**.

    e. Ensure **Change compute resource only** is selected.

    f. Click **Next**.

    g. Select a suitable compute host from the two hosts not entering maintenance mode.

    h. Click **Next** to accept defaults for all but the last dialog screen.

    i. Click **Finish** to migrate the VM.

11. When the host is in maintenance mode, complete the maintenance activity.

    Note: For a three-host vSAN cluster running ENM, VMs become non-compliant with the vSAN storage policy FTT=1 as all three hosts are required to enforce the policy. For this reason keep maintenance activities as short as possible.

Stop!

If it is necessary to reboot or power off the host in maintenance mode, perform a graceful shutdown by executing the appropriate command on the ESXi host.

```
# esxcli system shutdown reboot -r maintenance
```

or

```
# esxcli system shutdown poweroff -r maintenance
```

12. Exit maintenance mode.

    a. Expand the cluster inventory object `<vcenter_object_prefix>_CLUS` in **Navigator** to display the ESXi hosts.

    b. Right-click the host in maintenance mode and select **Maintenance Mode** > **Exit Maintenance Mode**.

    Note:
    — Progress appears in **Recent Tasks**.
    — vCenter removes the text `(maintenance mode)` next to the host in **Navigator** when maintenance mode is exited.
    — Running VMs regain compliance with the vSAN storage policy over time.

Note:
— The time needed to sync up on the returned host depends on the maintenance mode duration. If VMs are running, monitor their compliance after exiting maintenance mode to ensure they regain protection from the loss of an ESXi host. See section Confirm VM Storage Policy Compliance with vSAN Storage Policy on page 28.
— You can skip the remaining steps if they are being performed as part of an initial installation procedure.

13. Start the two instances that were gracefully shut down in *step 7*.

    a. Log on to VMS.

    b. List the servicereg and neo4j instances.

    ```
    [root@vms ~]# openstack server list -c Name -c Status |egrep -i "se →
    rvicereg|neo4j" | grep SHUTOFF
    ```

    c. Start the servicereq and neo4j instances in `SHUTOFF` state.

```
[root@vms ~]# openstack server start servicereg-X  <deployment-id>-neo4j-Y
```

d. Wait for the instance to reach status `Active`.

e. Log on to `servicereg-X` VM and as user `root` and end the consul process with the following command . This initiates a HA workflow for the VM.

```
# kill -15 $(pidof consul)
```

Example

```
[root@vms ~]# openstack server list |grep -i servicereg-X
| c3238857-3e8b-4740-a419-958ccdb78792 | servicereg-X | ACTIVE | vio_internal_network=<servicereg-X IP>
[root@vms ~]# ssh -i /vol1/senm/etc/<key_pair>.pem cloud-user@<servicereg-X IP>
[cloud-user@servicereg-X ~]$ sudo -i
[root@servicereg-X ~]# kill -15 $(pidof consul)
```

f. Log on to `<deployment_id>-neo4j-Y` VM and as user `root` and end the consul process with the following command.

```
# kill -15 $(pidof consul)
```

Note: This initiates a HA workflow for the VM.

14. Set vSphere DRS Automation Level to **Fully Automated** if it was changed in *step 10*.

15. Re-enable DRS rule `ControlStackAffinityRule` if it was disabled in step 9.

16. Follow the section, Confirm VM Storage Policy Compliance with vSAN Storage Policy on page 28.

17. Enable scheduled backups.

    — On deployments that use OMBS as the backup and restore storage solution: Follow the section *Reactivate Small Integrated ENM OMBS Policy* in Small Integrated ENM Backup and Restore System Administrator Guide .

    — On deployments that use the customer-provided NFS Share as the backup and restore storage solution: Follow the section *Reactivate Small Integrated ENM Backup Policy to External Device or Customer-Provided NFS Share* in Small Integrated ENM Backup and Restore System Administrator Guide.

Results

— Maintenance mode was achieved.

— Maintenance activity was completed.

— Host exited maintenance mode and all vSAN storage objects are compliant with the vSAN storage policy.

— Instances that were gracefully shut down from vCenter to enter maintenance mode were restarted from the Horizon dashboard and have status `Active`.

10.2 Planned Platform Shutdown Procedure

> This section describes steps to perform a planned platform shutdown of Small Integrated ENM Multi-Technology.

The procedure shuts down ENM using a VNF-LCM workflow before shutting down platform services. The ESXi hosts are placed in maintenance mode and powered off.

Note: Changes to the ENM and VNF-LCM SEDs between shutdown and startup are not supported by the procedure. Moving the deployment to a new site after shutdown is only possible if the same infrastructure configuration is in place at the new site.

Prerequisites

A full system backup is available.

Steps

1. Perform a platform health check by following the section Perform Platform Health Check on page 13.

   Note: Before the shutdown, the health check confirms all VMs and attached volumes are compliant with the vSAN storage policy and the OpenStack deployment has status `Running`.

2. Disable scheduled backups.

   — On deployments that use OMBS as the backup and restore storage solution: Follow the section *Deactivate Small Integrated ENM OMBS Policy* in Small Integrated ENM Backup and Restore System Administrator Guide

   — On deployments that use the customer-provided NFS Share as the backup and restore storage solution: Follow the section *Deactivate Small Integrated ENM Backup Policy to External Device or Customer-Provided NFS Share* in Small Integrated ENM Backup and Restore System Administrator Guide .

3. Allow active VNF-LCM workflows to complete.

a. Open the VNF-LCM UI `http://<external_ipv4_for_services_vm>/index.html#workflows` in a browser.

   Note: The value for `<external_ipv4_for_services_vm>` is available in the VNF-LCM SED.

b. Check for running workflows and decide if they can be canceled or allowed to complete before the shutdown.

4. Log on to the VMS on the management network **<vms_ip_vio_mgt>** as **root** user with the password **<vms_root_password>**, where `vms_ip_vio_mgt` and `vms_root_password` are parameters in the ENM SED.

5. Stop the DDC service.

```
# systemctl stop ddc.service
```

6. Stop the autofs service.

```
# systemctl stop autofs.service
```

7. Shut down ENM and VNF-LCM services from the VMS (a graceful shutdown is recommended and is the default mode). Follow Shut down ENM on page 79. Wait for the `manage_enm` script to complete on the VMS.

   Note: The VNF-LCM workflow **Manage ENM - Shutdown** can be tracked from the VNF-LCM UI at browser URL `http://<external_ipv4_for_services_vm>/index.html#workflows`.

8. Confirm that only VNF-LCM servers exist and have status `SHUTOFF` on VMS as user `root`.

   Example

```
[root@vms ~]# openstack server list
+--------------------------------------------------------------------------------+
+--------------------------------------------------------------------------------+
---------------

ID    Name    Status    Networks    Image Name
+--------------------------------------------------------------------------------+
+--------------------------------------------------------------------------------+
---------------

9a4b625c-a1b5-469e-8f87-cb6e4461cd88    <deployment_id>_-vnflaf-services-0    |
  SHUTOFF    vio_internal_network_5584=10.10.0.6;    ERICrhelvnflafimage_CXP |
9032490-4.4.21
                vio_external_network_5584=141.137.244.5
75f673dc-0236-4b11-9519-dc0cbb284539    <deployment_id>-vnflaf-db-0    SHUTO |
FF    vio_internal_network_5584=10.10.0.7;    ERICrhelpostgresimage_CXP90324 |
91-3.4.21
                vio_external_network_5584=141.137.244.50
+--------------------------------------------------------------------------------+
+--------------------------------------------------------------------------------+
---------------
[root@vms ~]#
```

9. Shut down VIO.

a. From VMS, log on to OMS as `viouser`.

```
[root@vms]# ssh viouser@oms
[viouser@oms]#
```

b. Stop the OpenStack deployment.

```
[viouser@oms]# sudo viocli deployment stop
```

10. Shut down the VIO OMS vApp from the vCenter client.

    a. Select **Menu > Hosts and Clusters** in the vCenter client.

    b. Right-click the vApp containing VM `management-server` in **Navigator** and select **Power > Shut Down**.

    c. Wait for the OMS vApp to shut down.

11. Shut down the VMS from the vCenter client.

    a. Select **Menu > Hosts and Clusters** in vCenter client.

    b. Right-click the VMS in **Navigator** and select **Power > Shut Down Guest OS**.

    c. Wait for the VMS to shut down.

12. Shut down the vCSA from the ESXi Host Client.

    a. Select **Menu > Hosts and Clusters** in vCenter client.

    b. Select VM `<vcenter_object_prefix>_VCSA` in **Navigator**.

    c. Select the **Summary** tab and note the IP address of the ESXi host running the vCSA VM.

    d. Log on to the ESXi host client running the vCSA VM with browser URL `https://<esxi_hostX_ip_vio_mgt>/ui` as user `root` and password `<esxi_hostX_mgt_password>`.

    Note: Where `<esxi_hostX_ip_vio_mgt>` is the management IP address of the ESXi host running the vCSA VM noted above.

    e. Select **Navigator > Virtual Machines**.

    f. Right-click VM `<vcenter_object_prefix>_VCSA` and select **Guest OS > Shut down**.

    g. Wait for the vCSA VM to shut down.

    h. Write down the IP address of the ESXi host that was running the vCSA VM.

Note: The platform startup procedure uses the IP address to access the host client and start the VM.

13. Place all three ESXi hosts in maintenance mode and power off.

    a. Log on to the first ESXi hosts with IP `<esxi_host1_ip_vio_mgt>` as user `root` with password `<esxi_host1_mgt_password>`.

    b. Enter maintenance mode.

    ```
    # esxcli system maintenanceMode set -e true -m noAction
    ```

    c. Confirm that maintenance mode is enabled.

    ```
    # esxcli system maintenanceMode get
    ```

    d. Power off the ESXi host.

    ```
    # esxcli system shutdown poweroff -r Scheduled
    ```

    e. Repeat the steps for the other two ESXi hosts.

10.3 Planned Platform Startup Procedure

> This section describes steps to start up the Small Integrated ENM Multi-Technology platform after a planned shutdown.

The procedure starts platform services and confirms platform health before running a VNF-LCM workflow to start ENM.

Note: If the procedure fails and the issue cannot be resolved, a full system or an ENM-only restore must be performed, depending on where the failure occurred. ENM-only restore is selected if the platform is passing health checks. See:

- *Backup and Restore with OMBS* in Small Integrated ENM Backup and Restore System Administrator Guide for ENM deployments that use OMBS as the backup and restore solution.
- *Backup and Restore with Customer-Provided NFS Share* in Small Integrated ENM Backup and Restore System Administrator Guide for ENM deployments that use the customer-provided NFS Share as the backup and restore solution.

Prerequisites

- The system was shut down following Planned Platform Shutdown Procedure on page 38.

Note: This procedure cannot be used to recover the platform. For recovery see Recover from a Complete Outage of Small Integrated ENM Multi-Technology on page 54.

— The ENM and VNF-LCM SEDs have not changed since the planned shutdown. Moving the deployment to a new site before startup is only possible if the same infrastructure configuration is in place at the new site.

— A system backup taken before the planned shutdown is available.

Steps

1. Power on all three ESXi hosts.

   Note: The hosts can be powered on in any order, a few seconds apart.

   a. Log on to the iLO of each ESXi host on IP `<esxi_hostX_ip_ilo>` as user `<esxi_hostX_ilo_user>` with password `<esxi_hostX_ilo_password>`.

   b. Select **Power Switch > Momentary Press**.

   c. Wait for all three hosts to show the **ESXi DCUI** console display.

2. Exit maintenance mode on all three ESXi hosts.

   a. Log on to the first ESXi hosts with IP `<esxi_host1_ip_vio_mgt>` as user `root` with password `<esxi_host1_mgt_password>`.

   b. Exit maintenance mode using the following command.

   ```
   # esxcli system maintenanceMode set -e false
   ```

   c. Confirm maintenance mode is disabled.

   ```
   # esxcli system maintenanceMode get
   ```

   d. Repeat the steps on the other two ESXi hosts.

3. Start up the `vCSA` VM.

   a. Log on to the ESXI host client with browser URL `https://<esxi_hostX_ip_vio_mgt>/ui` as user `root` with password `<esxi_hostX_mgt_password>`.

   Note:
   — Where `<esxi_hostX_ip_vio_mgt>` is the management IP address of the ESXi host that was running the `vCSA` VM noted from the platform shutdown procedure.

   — If it is not known which of the ESXi hosts was running `vCSA`, log on to each ESXi host client and use **Navigator** to locate VM `<vcenter_object_prefix>_VCSA`

b. Select **Navigator > Virtual Machines**.

c. Right-click VM `<vcenter_object_prefix>_VCSA` and select **Power > Power on**.

d. Wait for the vCenter client to start accepting logins at URL `https://<vcenter_ip_vio_mgt>/ui/`.

4. Start up the VMS from the vCenter client.

   a. Log on to the vCenter client at `https://<vcenter_ip_vio_mgt>/ui/` as user `administrator@vsphere.local` with password `<vcenter_sso_password>` from the SED.

   b. Select **Menu > Hosts and Clusters** in the vCenter client.

   c. Right-click the VMS in **Navigator** and select **Power > Power On**.

   d. Open a VM console and wait for the VMS to start up.

5. Start up the VIO OMS vApp from the vCenter client.

   a. Select **Menu > Hosts and Clusters** in the vCenter client.

   b. Right-click the vApp containing VM `management-server` in **Navigator** and select **Power > Power On**.

   c. Open a VM console and wait for the `management-server` VM to start up.

6. Start up the OpenStack deployment.

   Note: If the OpenStack deployment is not visible in vCenter, follow VIO Deployment Not Visible in vCenter GUI on page 132.

   a. Log on to the VMS as `root` user.

   b. Log on to OMS server as `root` user.

   ```
   [root@vms ~]# ssh viouser@oms
   viouser@oms:~# sudo -i
   ```

   c. Start all OpenStack services.

   ```
   root@oms:~# viocli services start
   ```

   d. Wait for the OpenStack deployment to reach the status `Running`.

7. Perform a platform health check (see section Perform Platform Health Check on page 13).

8. Start up ENM and VNF-LCM services from VMS; follow Start ENM on page 81 and wait for the `manage_enm` script to complete on VMS.

   Note:
   - The `manage_enm` script starts VNF-LCM services before running the `Manage ENM - Start` workflow. When VNF-LCM services are available, the workflow can be tracked from the VNF-LCM UI at URL `http://<external_ipv4_for_services_vm>/index.html#workflows`
   - If VNF-LCM service fails to start during `manage_enm`, log on to vCenter and open a VM console to VM `<deployment_id>-vnflaf-db-0`. If the VM tried to PXE boot instead of booting from the internal boot disk, follow section Recover VM in PXE Boot Mode on page 117. When the VM startup is complete, rerun the `manage_enm` script.
   - If the `Manage ENM - Start` workflow fails during `manage_enm`, recover by following section Recover ENM on page 85.

9. Follow ENM Health Check on page 22.

10. Enable scheduled ENM backups.
    - On deployments that use OMBS as the backup and restore storage solution: Follow the section *Reactivate Small Integrated ENM OMBS Policy* in Small Integrated ENM Backup and Restore System Administrator Guide.
    - On deployments that use the customer-provided NFS Share as the backup and restore storage solution: Follow the section *Reactivate Small Integrated ENM Backup Policy to External Device or Customer-Provided NFS Share* in Small Integrated ENM Backup and Restore System Administrator Guide.

10.4 Unplanned Loss of ESXI Host

This section details the steps to take in the unplanned loss of an ESXi host.

Note: This section applies to Small Integrated ENM Multi-Technology deployments only.

Small Integrated ENM Multi-Technology can tolerate the loss of a single ESXi host.

FM and vCenter alarms alert the user when an ESXi host fails. ENM HA workflows rebuild lost application VMs on the remaining two ESXi hosts.

ENM HA uses the Control Stack. If the lost ESXi host was running the Control Stack, ENM HA workflows start after the Control Stack VMs are restarted on another host.

The Control Stack restart takes about 20 minutes.

vSphere DRS rules exist to keep single-instance ENM applications off the ESXi host running the Control Stack so they can be recovered by ENM HA without waiting for a Control Stack restart.

Control Stack VMs are:

— `<vcenter_object_prefix>_VCSA` (vCenter Server Appliance)

— `management-server` (OpenStack Management Server)

— `OpenStack-ControlPlane-0` (VIO Controller)

— `OpenStack-Compute-0` (VIO Compute Driver)

— VMS (Virtual Management Server)

— `<deployment_id>`-vnflaf-db-0

— `<deployment_id>`-vnflaf-services-0

10.4.1 Recover after Unplanned Loss of ESXI Host

> This procedure describes how to track the recovery after losing an ESXi host and how to deal with issues that can occur.

Steps

1. Disable scheduled ENM backups.
   — On deployments that use OMBS as the backup and restore storage solution: Follow the section *Deactivate Small Integrated ENM OMBS Policy* in Small Integrated ENM Backup and Restore System Administrator Guide

   — On deployments that use the customer-provided NFS Share as the backup and restore storage solution: Follow the section *Deactivate Small Integrated ENM Backup Policy to External Device or Customer-Provided NFS Share* in Small Integrated ENM Backup and Restore System Administrator Guide .

2. Disable cron jobs that maintain the single-instance anti-affinity DRS rules.
   a. Edit the `root` cron table and comment out lines that call script `update_drs_rules.sh`.

      Note: Run on VMS as user `root`.

```
[root@vms]# crontab -e
#Ansible: drs_rule
#*/30 * * * * /opt/ericsson/senm/bin/update_drs_rules.sh -a update -c /vol1/senm/etc/drs_config.json
#Ansible: drs_cleanup
#45 4 * * * /opt/ericsson/senm/bin/update_drs_rules.sh -a cleanup -c /vol1/senm/etc/drs_config.json
*/30 * * * * /usr/bin/python /opt/ericsson/senm/bin/sienm_hc.py
#Ansible: backup_rule
1 0 * * 0 /opt/ericsson/bur-config-ombs/bin/senmbackup
```

b. Save the modified cron table.

3. Disable single-instance anti-affinity DRS rules.

   Note: Run on VMS as user `root`.

```
[root@vms]# /opt/ericsson/senm/bin/update_drs_rules.sh -a undo -c /vol1/senm/etc/drs_config.json
INFO: Starting new HTTP connection (1): haproxy-int
INFO: Starting new HTTP connection (1): haproxy-int
INFO: Found a rule for SISG with [ieatvio5557-vnflaf-services-0 (b516bc25-ffb8-4cce-b475-2fe9ff3ff6ff)] as member
INFO: Task finished successfully
INFO: Found a rule for SISG with [ieatvio5557-vnflaf-services-0 (b516bc25-ffb8-4cce-b475-2fe9ff3ff6ff)] as member
INFO: Task finished successfully
..
```

4. Ensure that an ENM HA workflow starts to rebuild lost application VMs.

   a. Open the VNF-LCM services URL `http://<external_ipv4_for_services_vm>/index.html#workflows` in a browser.

      Note: If the VNF-LCM services URL fails to load, see section Recover VNF-LCM Services after a Control Stack Failover on page 47.

   b. Check for a running **High Availability Workflow**.

   c. Open the workflow and select tab **Workflow Log**.

   d. Expand the first **INFO** log to see the list of application VMs handled by the workflow.

   e. Monitor the HA workflow until all application VMs are recovered apart from `<deployment_id>-neo4j-X` and `servicereg-X`.

Note:
- Application VMs `<deployment_id>-neo4j-X` and `servicereg-X` can only be recovered after the lost ESXi host has returned as they are members of a triple instance anti-affinity server group.
- If a HA workflow fails for any other application VMs, see section Resolve ENM HA Workflow Issues after a Control Stack Failover on page 48.
- If the ESXi host has already returned, see section Return of Lost ESXi Host on page 51.
- If the ESXi host has not yet returned, the HA workflow continues to retry their recovery for up to 21 days.

5. Run the steps in section Return of Lost ESXi Host on page 51 when the ESXi host has rejoined the cluster.

10.4.1.1 Recover VNF-LCM Services after a Control Stack Failover

> This section describes steps to handle cases where Control Stack VNF-LCM service availability is delayed.

If an ESXi host fails and the host was running the Control Stack, vSphere HA restarts the Control Stack VMs on another host. When Control Stack services are available, an ENM HA workflow starts to recover lost application instances.

Steps

1. Check if the VNF-LCM services restart is still in progress.

   Note:
   - The single-instance DRS rules can sometimes block the vSphere HA restart of Control Stack VM `<deployment_id>-vnflaf-services-0` causing the VM to remain powered off in the vCenter inventory.
     - If this happens vSphere HA restarts VM `<deployment_id>-vnflaf-services-0` shortly after the single-instance DRS rules are disabled at the start of section Recover after Unplanned Loss of ESXI Host on page 45.
     - Check for VNF-LCM services URL to become available at `http://<external_ipv4_for_services_vm>/index.html#workflows`.

2. Restart VNF-LCM servers only if the services URL fails to load or it is not possible to SSH to the VNF-LCM servers after 20 minutes.

   a. Restart VNF-LCM servers.

Note: Run on VMS as user `root`.

```
[root@vms]# openstack server reboot <deployment_id>-vnflaf-db-0
[root@vms]# openstack server reboot <deployment_id>-vnflaf-services
-0
```

Note:
— If a VNF-LCM server fails to reboot, check if the VM is powered on in the vCenter inventory. If necessary, select and power on the VM from the vCenter inventory and then run the commands below on the VMS as user `root`. In this example, both of the above VNF-LCM servers failed to reboot.

```
# Set ACTIVE state
[root@vms]# openstack server set --state active <deployment_id>-vnflaf-db-0
[root@vms]# openstack server set --state active <deployment_id>-vnflaf-services-0

# confirm ACTIVE state
[root@vms]# openstack server list | grep -i vnf

# Restart
[root@vms]# openstack server stop <deployment_id>-vnflaf-services-0
[root@vms]# openstack server stop <deployment_id>-vnflaf-db-0
[root@vms]# openstack server start <deployment_id>-vnflaf-db-0
[root@vms]# openstack server start <deployment_id>-vnflaf-services-0

# confirm ACTIVE state
[root@vms]# openstack server list | grep -i vnf
```

— If server `<deployment_id>-vnflaf-db-0` still fails to reboot, use the OpenStack Horizon Web UI to open a server console. If error `Operating System not found` is displayed, apply the boot recovery described in section Recover VM in PXE Boot Mode on page 117.

3. Confirm that an ENM HA workflow starts to recover lost instances after VNF-LCM services are available at URL `http://<external_ipv4_for_services_vm>/index.html#workflows`

10.4.1.2 Resolve ENM HA Workflow Issues after a Control Stack Failover

> This section shows how to resolve ENM HA workflow issues after a control stack failover.

The ENM HA workflow that starts after losing an ESXi host rebuilds all lost application VMs apart from `<deployment_id>-neo4j-X` and `servicereg-X`. If the lost ESXi host was running the Control Stack, vSphere HA must first restart the Control Stack VMs before ENM HA workflows can start.

Run the following steps if the ENM HA workflow fails to rebuild application VMs other than `<deployment_id>-neo4j-X` and `servicereg-X`.

Note: Check the list of successfully recovered VMs in the ENM HA workflow `INFO` messages that follow an initial `ERROR` message. The `ERROR` message may show that the workflow could not recover VMs other than `<deployment_id>-neo4j-X` and `servicereg-X` but a later retry may have succeeded.

Steps

Check that OpenStack Services are up.

1. Ensure OpenStack cinder services are up.

    a. Connect over SSH to VMS as user `root`.

```
[root@vms]# openstack volume service list
+------------------+-----------------------------------+------+----  →
-----+-------+----------------------------+
| Binary           | Host                              | Zone | Sta  →
tus  | State | Updated At                 |
+------------------+-----------------------------------+------+----  →
-----+-------+----------------------------+
| cinder-backup    | loadbalancer01                    | nova | ena  →
bled | up    | 2019-10-24T09:32:45.000000 |
| cinder-volume    | loadbalancer01@nova:10.149.58.163 | nova | ena  →
bled | down  | 2019-10-24T09:32:36.000000 |
| cinder-scheduler | loadbalancer01                    | nova | ena  →
bled | up    | 2019-10-24T09:32:42.000000 |
+------------------+-----------------------------------+------+----  →
-----+-------+----------------------------+
[root@vms]#
```

    b. Connect over SSH to `loadbalancer01` as `root` if a service is down.

```
[root@vms]# ssh viouser@oms
[viouser@oms]$ ssh loadbalancer01
[viouser@loadbalancer01]$ sudo -i

# Restart service
[root@loadbalancer01]# systemctl restart cinder-volume

# Ensure service is enabled
[root@loadbalancer01]# systemctl enable cinder-volume and
                               restart the service as user
```

    c. Ensure the OpenStack Glance API service is up on `loadbalancer01`.

```
# Check glance-api service
[root@loadbalancer01]# systemctl | grep glance

# Restart glance-api service if not running
[root@loadbalancer01]# systemctl restart glance-api

# Ensure glance-api service is enabled
[root@loadbalancer01]# systemctl enable glance-api
```

d. Ensure the OpenStack Neutron Server service is up on `loadbalancer01`.

```
# Check neutron-server service
[root@loadbalancer01]# systemctl | grep neutron

# Restart neutron-server service if not running
[root@loadbalancer01]# systemctl restart neutron-server

# Ensure neutron-server service is enabled
[root@loadbalancer01]# systemctl enable neutron-server
```

2. Remove stale ENM volume attachments.

   Note: Run on VMS as user `root`.

   a. Check for volumes that are attached to a stale server ID instead of a server name.

```
[root@vms]# openstack volume list
.
| 8100572c-b2ff-417f-a445-6953a3a6e86c | vio-5646-ops_volume-1       →
          | in-use | 1  | Attached to b634c282-975a-421a-bd41-12131  →
a289f2d on /dev/sdb |
| 4f5f76d0-a48a-4425-973c-22f03436309f | vio-5646-neo4j_volume-2     →
          | in-use | 60 | Attached to b4b002b3-674e-4340-994c-84dc6  →
0e31a39 on /dev/sdb |
| 3f91027c-2d87-49f2-af15-d36ee90cd1c2 | vio-5646-serviceregistry_v  →
olume-1   | in-use | 1  | Attached to 07b76af6-3241-43c3-b9dc-30fd8  →
27be405 on /dev/sdb |
..
```

   b. Set state `detached` and `available` for volumes attached to stale server IDs.

```
[root@vms]# cinder reset-state --attach-status detached --state ava  →
ilable 8100572c-b2ff-417f-a445-6953a3a6e86c
[root@vms]# cinder reset-state --attach-status detached --state ava  →
ilable 4f5f76d0-a48a-4425-973c-22f03436309f
[root@vms]# cinder reset-state --attach-status detached --state ava  →
ilable 3f91027c-2d87-49f2-af15-d36ee90cd1c2
```

3. Reset volume state for detached volumes stuck in state `reserved`.

   Note:
   - If the ENM HA workflow fails to rebuild an application VM, check if it has a detached volume with state `reserved`.
   - If the lost ESXi host has not yet returned, ignore volumes in state `reserved` for VMs `<deployment_id>-neo4j-X` and `servicereg-X`. These are handled in section Return of Lost ESXi Host on page 51.

   a. Check for a detached volume in state `reserved`.

      Note: Run on VMS as user `root`.

```
[root@vms]# openstack volume list | grep -i reserved
| e708df26-0c3c-4adf-bec4-ec7531cda199 | vio-5646-ops_volume-1 | re →
served | 1 | |
```

b. Confirm the ENM HA workflow is failing to attach the volume on the VNF-LCM services server.

```
[root@vms]# ssh -i /vol1/senm/etc/<key_name>.pem cloud-user@vnflaf- →
services
[cloud-user@vio-5646-vnflaf-services-0 ~]$ sudo grep -i "status mus →
t be available" /ericsson/3pp/jboss/standalone/log/server.log

HA_1571761575428 ... Invalid volume: Invalid input received: Invali →
d volume: Volume e708df26-0c3c-4adf-bec4-ec7531cda199 status must b →
e available or downloading (HTTP 400)
```

c. Reset the volume state to `detached` and `available`.

```
[root@vms]# cinder reset-state --attach-status detached --state ava →
ilable e708df26-0c3c-4adf-bec4-ec7531cda199
```

4. Confirm that the next retry of the ENM HA workflow rebuilds the application VM.

10.4.2 Return of Lost ESXi Host

> This section details the steps to follow for a returning ESXi host.

Note: This section applies to Small Integrated ENM Multi-Technology only.

When the lost ESXi host is powered up again and rejoins the cluster, vSAN storage objects on the host are synchronized to enforce the vSAN storage policy FTT=1 that protects against loss of an ESXi host.

vSphere DRS starts to migrate VMs onto the returned host.

Steps

1. Monitor DRS VM migrations onto the returned host in vCenter **Recent Tasks** to confirm that the host is accepting workloads.
2. Enable cron jobs for single-instance anti-affinity DRS rules.

   a. Edit the root cron table and activate lines that call script `update_drs_rules.sh`.

   Note: Run on VMS as user `root`.

```
[root@vms]# crontab -e
```

```
#Ansible: drs_rule
*/30 * * * * /opt/ericsson/senm/bin/update_drs_rules.sh -a update -c /vol1/senm/etc/drs_config.json
#Ansible: drs_cleanup
45 4 * * * /opt/ericsson/senm/bin/update_drs_rules.sh -a cleanup -c /vol1/senm/etc/drs_config.json
*/30 * * * * /usr/bin/python /opt/ericsson/senm/bin/sienm_hc.py
#Ansible: backup_rule
1 0 * * 0 /opt/ericsson/bur-config-ombs/bin/senmbacku
```

b. Save the modified cron table.

3. Enable single-instance anti-affinity DRS rules.

   Note: Run on VMS as user `root`.

```
[root@vms]# /opt/ericsson/senm/bin/update_drs_rules.sh -a update -c /vol1/senm/etc/drs_config.json
INFO: Starting new HTTP connection (1): haproxy-int
INFO: Starting new HTTP connection (1): haproxy-int
INFO: Add [ieatvio5557-vnflaf-services-0 (b516bc25-ffb8-4cce-b475-2fe9ff3ff6ff)] to [ieatvio5557-cnom-0 (939f7327-2649-4e5c-9a4d-5803bf07692a)]'s rule
INFO: Task finished successfully
INFO: Add [ieatvio5557-vnflaf-services-0 (b516bc25-ffb8-4cce-b475-2fe9ff3ff6ff)] to [ieatvio5557-domainproxy-0 (f758d0e2-2db5-4298-8e75-48010a8f4712)]'s rule
INFO: Task finished successfully
..
```

4. List and delete OpenStack servers with status `ERROR` for lost instance `<deployment_id>-neo4j-X`.

   Note: Run on VMS as user `root`.

```
# List
[root@vms]# openstack server list | grep neo4j | grep ERROR
| 16bc2870-ea47-4b46-83ba-6579970c0f5f | vio-5646-neo4j-0 | ERROR |
| cffa3f27-b92e-4b0e-b014-43eb4b1c8d2b | vio-5646-neo4j-0 | ERROR |
| b036170d-b885-498e-a007-386c2c16b6aa | vio-5646-neo4j-0 | ERROR |
..

# Delete
[root@vms]# openstack server delete 16bc2870-ea47-4b46-83ba-6579970c0f5f \
> cffa3f27-b92e-4b0e-b014-43eb4b1c8d2b \
> b036170d-b885-498e-a007-386c2c16b6aa \
..
```

5. List and reset the volume state for lost instance `<deployment_id>-neo4j-X` if the volume state is `reserved`.

   Note: Run on VMS as user `root`.

```
# List
[root@vms]# openstack volume list | grep neo4j
| 4f5f76d0-a48a-4425-973c-22f03436309f | vio-5646-neo4j_volume-2 | in-use
| 60 | Attached to vio-5646-neo4j-2 on /dev/sdb |
| f304ddfc-5fd4-4af8-bf4a-42741b38455f | vio-5646-neo4j_volume-1 | in-use
| 60 | Attached to vio-5646-neo4j-1 on /dev/sdb |
| 2f48b006-cffa-4460-b5f2-b027d633a09c | vio-5646-neo4j_volume-0 | reserved
| 60 |                                      |

# Reset
```

```
[root@vms]# cinder reset-state --attach-status detached --state available 2f
48b006-cffa-4460-b5f2-b027d633a09c
```

6. List and delete OpenStack servers with status `ERROR` for lost instance `servicereg-X`.

   Note: Run on VMS as user `root`.

```
# List
[root@vms]# openstack server list | grep servicereg | grep ERROR
| 1c43da2b-197b-4090-bdaf-f7818ba8a838 | servicereg-2 | ERROR |
| 3ca6b97f-37a3-4926-ab38-83715845f6af | servicereg-2 | ERROR |
| c0cfc68a-8fff-40b4-a810-4deaf6f56497 | servicereg-2 | ERROR |
..

# Delete
[root@vms]# openstack server delete 1c43da2b-197b-4090-bdaf-f7818ba8a838 \
> 3ca6b97f-37a3-4926-ab38-83715845f6af \
> c0cfc68a-8fff-40b4-a810-4deaf6f56497 \
..
```

7. List and reset the volume state for lost instance `servicereg-X` if the volume state is `reserved`.

   Note: Run on VMS as user `root`.

```
# List
[root@vms]# openstack volume list | grep servicereg

| e708df26-0c3c-4adf-bec4-ec7531cda199 | vio-5646-serviceregistry_volume-0 |
 in-use   | 1 | Attached to servicereg-1 on /dev/sdb |
| 9ece7890-14e7-4b6a-a63d-fc5fa91d48df | vio-5646-serviceregistry_volume-2 |
 in-use   | 1 | Attached to servicereg-3 on /dev/sdb |
| 3f91027c-2d87-49f2-af15-d36ee90cd1c2 | vio-5646-serviceregistry_volume-1 |
 reserved | 1 |                                      |

# Reset
[root@vms]# cinder reset-state --attach-status detached --state available 3f
91027c-2d87-49f2-af15-d36ee90cd1c2
```

8. Monitor the ENM HA workflow that is still running if the lost ESXi host was returned within 21 days.

   Note:
   - The workflow can now complete as it has a third host available for the third instance of `neo4j` and `servicereg`.
   - See section Manual Recovery of an ENM Server Instance on page 25 if the HA workflow terminated before the ESXi host returned.

9. Reconfigure the Control Stack restart condition in vCenter cluster VM overrides. Log on to VMS and run the following command.

```
[root@vms]# ansible-playbook /opt/ericsson/edp/automation/sienm/playbooks/si
enm_configure_control_stack.yml
```

10. Enable scheduled ENM backups.

— On deployments that use OMBS as the backup and restore storage solution: Follow the section *Reactivate Small Integrated ENM OMBS Policy* in Small Integrated ENM Backup and Restore System Administrator Guide .

— On deployments that use the customer-provided NFS Share as the backup and restore storage solution: Follow the section *Reactivate Small Integrated ENM Backup Policy to External Device or Customer-Provided NFS Share* in Small Integrated ENM Backup and Restore System Administrator Guide.

10.5 Recover from a Complete Outage of Small Integrated ENM Multi-Technology

This section describes steps to recover a Small Integrated ENM Multi-Technology system after a power or network outage on two or all three ESXi hosts.

— *Backup and Restore with OMBS* in Small Integrated ENM Backup and Restore System Administrator GuideThis procedure is a best effort recovery. If this procedure fails, then a full system restore from backup must be performed, see: for ENM deployments that use OMBS as the backup and restore solution.

— *Backup and Restore with Customer-Provided NFS Share* in This procedure is a best effort recovery. If this procedure fails, then a full systemSmall Integrated ENM Backup and Restore System Administrator Guide for ENM deployments that use the customer-provided NFS Share as the backup and restore solution.

Note: If no backup is available, the system must be reinstalled.

Note:

— The system can tolerate the loss of a single host. For loss of a single host, refer to Unplanned Loss of ESXI Host on page 44.

— It may not be necessary to run all steps. For example, power may already have been restored to all three hosts before recovery is started.

Stop!

The full recovery procedure starts with all ESXi hosts powered off. If two hosts lost power, power off the remaining host before starting. Ensure all hosts are powered off before starting the recovery procedure.

When the outage is over, vSphere HA starts the platform Control Stack automatically. The Control Stack is the group of VMs listed below that provide OpenStack cloud and VNF-LCM workflow services. The Control Stack takes about 20 minutes to start.

System recovery can be performed by the `manage_enm` script in this procedure if the following conditions are met:

— The platform Control Stack is recovered after the outage.

— All ENM volumes are passing vSAN health checks. Steps are provided to identify the ENM volumes and their vSAN health.

The `manage_enm` script creates new ENM servers and attaches the existing volumes.

If power or network was already restored before starting recovery, the first step is to check if the vCenter client is available or starting up.

Steps are described to resolve any file system errors encountered during the vCSA boot phase.

Note: — vCenter client is used to confirm the health of the Control Stack.

— vCenter client is used to confirm vSAN cluster health and the health of ENM volumes.

The health of the VMS is checked and the two VNF-LCM servers are restarted from the VMS if necessary.

After a network outage, when the Control Stack is operational, check the VNF-LCM services web UI for running ENM HA workflow instances.

Note: Monitor the workflows. If there are only a few running, they may recover the ENM system without the need to run the `manage_enm` script. If the number of concurrent HA workflows increases or workflows appear to hang it is quicker to run the `manage_enm` script.

After the recovery some time is needed to sync with the radio network before all ENM features are fully available.

After a power outage to the cluster, VIO OpenStack may restart all ENM applications and they may appear healthy.

Note: It is still necessary to run `manage_enm` to complete the recovery as most applications are not designed to be restarted. They must be recreated from a VNF-LCM workflow triggered by the `manage_enm` script.

Control Stack VMs are:

Table 2 Control Stack VMs

VM Name	Description
<vcenter_object_prefix>_VCSA	vCSA appliance VM running vCenter.
VMS	Virtual Management Server. Runs deployment and configuration scripts.
management-server	OpenStack Management Server. Manages the VIO deployment.
Openstack-Compute-0	VIO Compute.
Openstack-ControlPlane-0	VIO Controller.
<deployment_id>-vnflaf-services-0	VNF-LCM. Lifecycle Management services.
<deployment_id>-vnflaf-db-0	VNF-LCM database server.

Note:
— If a Control Stack VM fails to boot because of file system errors displayed on the VM console, follow Resolve File System Errors on page 134.

— After a network outage, if a Control Stack VM is unresponsive, reboot the VM. ENM HA can attempt system recovery if the Control Stack is operational.

Prerequisites

— HPE iLO access to all three ESXi hosts.

— Site Engineering Document (SED) is available.

— No network or hardware faults exist.

— Two or all three ESXi hosts lost power or network connectivity.

Steps

1. Power on all three ESXi hosts.

    a. Log on to the iLO of each ESXi host with IP <esxi_hostX_ip_ilo> as user <esxi_hostX_ilo_user> with password <esxi_hostX_ilo_password>.

    b. Select **Power Switch > Momentary Press**.

    Note:
    — The hosts can be powered on in any order, a few seconds apart.

    — Ensure to power on all hosts within a few seconds of each other.

    c. Wait for all three hosts to show the ESXi DCUI console login screen.

Note: vSAN initialization takes longer than usual during the boot phase.

2. Check vSphere HA started the vCSA VM.

   a. Locate the ESXi host running the Control Stack by logging on to each ESXi host client with browser URL `https://<esxi_hostX_ip_vio_mgt>` as user `root` with password `<esxi_hostX_mgt_password>`.

   b. In **Navigator**, select **Virtual Machines** and search for VM `<vcenter_object_prefix>_VCSA`.

   c. Select and power on the vCSA VM if vSphere HA has not started the vCSA VM within five minutes.

   d. Open a console to the vCSA VM in the ESXi host client to monitor boot progress.

      Note: If a console message is displayed to run a manual `fsck` command, follow Resolve File System Errors on page 134.

   e. Wait for the vCenter Client to become available and accepting logins at browser URL `https://<vcenter_ip_vio_mgt>/ui/`.

   f. Log on to the vCenter Client as user `administrator@vsphere.local` with password `<vcenter_sso_password>`.

   g. Confirm all three ESXi hosts are connected to the vCenter cluster.

      Note: If any ESXi host is in maintenance mode, right-click the host in **Navigator** and select **Exit maintenance mode**.

3. Check vSphere HA started the VMS VM.

   a. Select and power on the VMS VM in vCSA **Navigator** if not already powered on.

   b. Open a console to the VM to monitor boot progress until accepting logins.

   c. Log on to VMS at `<vms_ip_vio_mgt>` as user `root` with password `<vms_root_password>`.

   d. Confirm the backup mount on the VMS if a backup and restore solution is configured.

```
[root@vms ~]# mount | grep -i backup
10.10.10.10:/vENM_backup_5653 on /BACKUP type nfs4 (rw,relatime,ver →
s=4.0,rsize=1048576,wsize=1048576,namlen=255,hard,proto=tcp,port=0, →
timeo=600,retrans=2,sec=sys,clientaddr=10.10.10.11,local_lock=none, →
addr=10.10.10.10)
```

4. Check vSphere HA started the OMS vApp.

a. Select and power on the vApp containing the OMS VM `management-server` in vCSA **Navigator** if not already powered on.

b. Open a console to the OMS VM to monitor boot progress until accepting logins.

5. Check vSphere HA started OpenStack Services.

a. Click the **Menu** icon and select **VMware Integrated OpenStack** in VCSA **Navigator**.

Note: If the OpenStack deployment is not visible in vCSA, follow Recover OMS connectivity to VIO on page 10.

b. Use vCenter to gracefully shut down VMs `Openstack-ControlPlane-0`, `Openstack-Compute-0`, and the vApp containing `management-server` if vCenter still can't access the OpenStack deployment and then start the vApp containing `management-server`.

Note: The OpenStack deployment is started automatically.

c. Select **OpenStack Deployments**.

d. Select **All Actions > Start OpenStack Services** if OpenStack Services are not started.

e. Wait for the OpenStack deployment to achieve status `Running`.

f. Ensure OpenStack cinder services are up by connecting through SSH to VMS as user `root`.

```
[root@vms]# openstack volume service list
+------------------+-------------------------------------+------+----  →
-----+-------+----------------------------+
| Binary           | Host                                | Zone | Sta  →
tus  | State | Updated At                 |
+------------------+-------------------------------------+------+----  →
-----+-------+----------------------------+
| cinder-backup    | loadbalancer01                      | nova | ena  →
bled | up    | 2019-10-24T09:32:45.000000 |
| cinder-volume    | loadbalancer01@nova:10.149.58.163 | nova | ena  →
bled | down  | 2019-10-24T09:32:36.000000 |
| cinder-scheduler | loadbalancer01                      | nova | ena  →
bled | up    | 2019-10-24T09:32:42.000000 |
+------------------+-------------------------------------+------+----  →
-----+-------+----------------------------+
[root@vms]#
```

g. Connect through SSH to `loadbalancer01` and restart the service as user `root` if a service is down.

```
[root@vms]# ssh viouser@oms
[viouser@oms]$ ssh loadbalancer01
[viouser@loadbalancer01]$ sudo -i

# Restart service
[root@loadbalancer01]# systemctl restart cinder-volume
```

```
# Ensure service is enabled
[root@loadbalancer01]# systemctl enable cinder-volume
```

h. Ensure OpenStack Glance API service is up on `loadbalancer01`.

```
# Check glance-api service
[root@loadbalancer01]# systemctl | grep glance

# Restart glance-api service if not running
[root@loadbalancer01]# systemctl restart glance-api

# Ensure glance-api service is enabled
[root@loadbalancer01]# systemctl enable glance-api
```

i. Ensure OpenStack Neutron Server service is up on `loadbalancer01`.

```
# Check neutron-server service
[root@loadbalancer01]# systemctl | grep neutron

# Restart neutron-server service if not running
[root@loadbalancer01]# systemctl restart neutron-server

# Ensure neutron-server service is enabled
[root@loadbalancer01]# systemctl enable neutron-server
```

6. Check vSAN cluster datastore health.

   a. Select the cluster inventory object `<vcenter_object_prefix>_CLUS` in vCSA **Navigator**.

   b. Select **Monitor > vSAN > Skyline Health**.

   c. Confirm that all health check have passed apart from the following:

      — **Hardware compatibility** tests

      — **Disk format version** (This test can be disabled by clicking **Disk format version** > **Silence Alert** > **Yes**)

      — **vSAN Build Recommendation** (This test can be disabled by clicking **vSAN Build Recommendation** > **Silence Alert** > **Yes**)

      — **vSAN Disk Balance**

   d. Click the **Retest** button to refresh the test result if a health check has failed.

Note:
- Critical checks are **Network**, **Physical Disk**, **Data**, **Cluster** and **Limits**.
- If **Data** health check **vSAN object health** has failed, click **Repair Objects immediately** followed by the **Retest** button.
- If **Cluster** health check **vCenter state is authoritive** has failed, click button **Update ESXi configuration** followed by the **Retest** button.

7. Confirm the health of all ENM volume objects on the vSAN if the vSAN **Data** health check shows inaccessible data objects.

   Note: ENM volumes names start with `<deployment_id>` and contain `vol` or `volume` in the name.

   a. Select **Virtual Objects** on the **Monitor > vSAN** tab.

   b. Click the filter icon beside name and enter **vol** in the filter box to list all ENM volumes.

   c. Confirm all volumes are displayed as `Healthy` under the **vSAN Object Health** column.

      Note: If any ENM volume is listed as `Inaccessible`, a full system restore from backup must be performed.

8. Restart the VNF-LCM servers if the VNF-LCM services URL fails to load at URL `http://<external_ipv4_for_services_vm>/index.html#workflows` .

   a. Run on VMS as user `root`.

```
[root@vms]# openstack server reboot <deployment_id>-vnflaf-db-0
[root@vms]# openstack server reboot <deployment_id>-vnflaf-services →
-0
```

   b. Check that the VM is powered on in the vCenter inventory if a VNF-LCM server fails to reboot (if necessary select and power on the VM from the vCenter inventory and then run the commands below on the VMS as user `root`).

      Note: In this example both VNF-LCM servers failed to reboot.

```
# Set ACTIVE state
[root@vms]# openstack server set --state active <deployment_id>-vnf →
laf-db-0
[root@vms]# openstack server set --state active <deployment_id>-vnf →
laf-services-0

# confirm ACTIVE state
[root@vms]# openstack server list | grep -i vnf
```

```
# Restart
[root@vms]# openstack server stop <deployment_id>-vnflaf-services-0
[root@vms]# openstack server stop <deployment_id>-vnflaf-db-0
[root@vms]# openstack server start <deployment_id>-vnflaf-db-0
[root@vms]# openstack server start <deployment_id>-vnflaf-services-
0

# confirm ACTIVE state
[root@vms]# openstack server list | grep -i vnf
```

Note: If server `<deployment_id>-vnflaf-db-0` still fails to reboot, use the OpenStack Horizon Web UI to open a server console. If error `Operating System not found` is displayed, apply the boot recovery described in section Recover VM in PXE Boot Mode on page 117.

9. Wait for the VNF-LCM services web UI to become available and accepting logins at browser URL `http://<external_ipv4_for_services_vm>/index.html#workflows`.

   Note: The value for `<external_ipv4_for_services_vm>` is available in the VNF-LCM SED.

10. For a network outage, check the VNF-LCM web interface for active ENM HA workflows.

    Note:
    - Monitor the workflows when running the platform health checks in the next step. If there are only a few running they may recover the ENM system without the need to run the `manage_enm` script.
    - If the number of concurrent ENM HA workflows increases or workflows appear to hang it is quicker to run the `manage_enm` script.

11. Ensure the `vapi` service is running on the OMS.

    a. SSH to the VMS on the management network IP `<vms_ip_vio_mgt>` as `root` with password `<vms_root_password>` from the ENM SED.

    b. SSH to the OMS from the VMS as user `viouser`; switch to `root` user; and restart the `vapi` service.

    ```
    [root@vms] # ssh viouser@oms
    [viouser@oms] # sudo -i
    [root@oms] # systemctl start vapi
    ```

12. Follow Perform Platform Health Check on page 13.

13. Disable scheduled backups.

    - On deployments that use OMBS as the backup and restore storage solution: Follow the section *Reactivate Small Integrated ENM OMBS Policy* in Small Integrated ENM Backup and Restore System Administrator Guide.

— On deployments that use the customer-provided NFS Share as the backup and restore storage solution: Follow the section *Reactivate Small Integrated ENM Backup Policy to External Device or Customer-Provided NFS Share* in Small Integrated ENM Backup and Restore System Administrator Guide.

14. Follow section Recover ENM on page 85.

    Note:
    — After a power outage to the cluster, VIO OpenStack may restart all ENM applications and they may appear healthy. It is still necessary to run `manage_enm` to complete the recovery as most applications are not designed to be restarted.

    — If the `manage_enm` script fails at the first attempt, re-run the script. If the second run fails contact customer support or restore the system from backup.

15. Reconfigure the Control Stack restart condition in vCenter cluster VM overrides. Log on to VMS and run the following command.

```
[root@vms]# ansible-playbook /opt/ericsson/edp/automation/sienm/playbooks/si →
enm_configure_control_stack.yml
```

16. Follow section ENM Health Check on page 22.

17. Enable scheduled backups.

    — On deployments that use OMBS as the backup and restore storage solution: Follow the section *Reactivate Small Integrated ENM OMBS Policy* in Small Integrated ENM Backup and Restore System Administrator Guide .

    — On deployments that use the customer-provided NFS Share as the backup and restore storage solution: Follow the section *Reactivate Small Integrated ENM Backup Policy to External Device or Customer-Provided NFS Share* in Small Integrated ENM Backup and Restore System Administrator Guide.

Results

The Small Integrated ENM Multi-Technology system is recovered and passing critical platform and ENM health checks.

10.6 Recover after Healthy Disk Goes Offline Because of Known Issue with HPE Smart Array Controller

> This section describes steps to recover from a known issue with HPE Smart Array controller firmware for Small Integrated ENM Multi-Technology deployments on Gen9 or Gen10 rack systems.

Note:
- Small Integrated ENM Transport Only is not affected.
- This recovery procedure only applies if the issue meets the criteria set out in the *Prerequisites*.

The issue causes a healthy vSAN physical disk to go offline, triggering a **vSAN All Paths Down (APD)** event with vCenter and FM alarms. The physical disk still displays as healthy on the HPE iLO **System Information** page of the associated ESXi host.

In vCenter, vSAN health checks fail under categories **Physical Disks** and **Data**. The affected VMs become non-compliant with the configured vSAN storage policy FTT=1.

When the one-hour timer delaying vSAN repair expires, the vSAN policy creates copies of the storage objects on the remaining disks and the VMs regain compliance. When compliance is reached, vSAN health checks begin to pass again under the category **Data**.

However, the vSAN physical disk remains unused with a status of **Absent** under the **Physical Disks** health check category which continues to fail.

When all VMs are compliant with the vSAN storage policy, the ESXi host can be placed in maintenance and rebooted to recover the absent disk and all vSAN health checks pass again.

Note: There is no ENM outage if all ESXi hosts in the cluster are fully functional and handling ENM workloads when the issue occurs.

Prerequisites

- All ESXi hosts in the cluster are running ENM workloads when the issue occurs.
- Some of, or all, the following alarms appear in vCenter and ENM FM.
  - vSAN health alarm `vSAN object health`
  - vSAN health alarm `Software state health`
  - vSAN health alarm `Metadata health`

  - vSAN health alarm `Overall disks health`
  - VIO Alarm Service: `<ESXi host IP address> Errors occurred on the disk(s) of a vSAN host.`

- Failed vSAN health checks under categories **Physical Disks** and **Data**.

  Note: To view the health checks in vCenter, select **<vcenter_object_prefix>_CLUS** > **Monitor** > **vSAN** > **Skyline Health**.

  The associated ESXi host is identified by expanding checks in the **Physical Disks** category. The disk with the issue is marked `Absent`.

  Note: vSAN health checks under the category **Data** pass if the issue was discovered after vSAN had time to reapply the storage policy to the affected VMs.

- If more than one disk is offline, they must all belong to the same ESXi host.
- The HPE iLO **System Information** page of the affected ESXi host does not indicate a storage device failure. This check distinguishes the event from an actual disk failure.

  Note: The health of each physical drive attachment under each storage controller has status `OK`.

- The ESXi host kernel log file `/scratch/log/vmkernel.log` contains entries similar to:

```
2019-06-07T10:22:34.342Z cpu32:7413659)smartpqi: pqi_AllocDmaFromMemPool:92: →
 vmk_MemPoolAlloc failed: Out of memory.
2019-06-07T10:22:35.362Z cpu0:65974)StorageApdHandlerEv: 110: Device or file →
system with identifier [mpx.vmhba0:C0:T69:L0] has entered the All Paths Down →
 state.
2019-06-07T10:22:37.372Z cpu31:67506)WARNING: LSOM: LSOMEventNotify:6956: Vi →
rtual SAN device 5247bf5b-afa8-62e9-f2dc-fb9ad2365a4c has gone offline.
```

Steps

1. Track vSAN health checks in vCenter until all health checks in the **Data** category have passed.

   a. Log on to the vCenter client URL `https://<vcenter_ip_vio_mgt>/ui/` as administrator user `administrator@vsphere.local` with password `<vcenter_sso_password>`.

   b. Select the cluster inventory object `<vcenter_object_prefix>_CLUS` in **Navigator**.

   c. Select **Monitor** > **vSAN** > **Skyline Health**.

   d. Expand **Data** and confirm check **vSAN object health** has passed.

2. Confirm all VMs are compliant with the vSAN storage policy (follow section Confirm VM Storage Policy Compliance with vSAN Storage Policy on page 28).

3. Place the ESXi host identified in the *Prerequisites* in maintenance mode (follow section ESXi Maintenance Mode with Healthy vSAN Cluster on page 31.

   Note: Stop before step Exit Maintenance Mode.

4. Right-click the ESXi host in maintenance mode in vCenter and select **Power** > **Reboot** to reboot the ESXi host and recover the absent disk.

5. Take the host out of maintenance mode when the ESXi host has returned (follow section ESXi Maintenance Mode with Healthy vSAN Cluster on page 31).

   Note: Proceed from step Exit Maintenance Mode.

6. Monitor vSAN health checks and confirm that they are all passing.

   a. Select **<vcenter_object_prefix>_CLUS** > **Monitor** > **vSAN** > **Skyline Health** in vCenter.

   b. Confirm that the disk is no longer marked **Absent** in the **Physical Disks** health check category.

11 Small Integrated ENM Transport Only Host Startup, Shutdown, and Recovery

This section describes the maintenance and recovery procedures for Small Integrated ENM Transport Only deployments

11.1 Shutdown Procedure for Small Integrated ENM Transport Only

This section provides steps to shut down a Small Integrated ENM Transport Only deployment gracefully.

Stop!

This procedure, along with Startup Procedure for Small Integrated ENM Transport Only on page 69, cannot be used in the following scenarios:

— Hardware replacement of some or all rack components has taken place.

— The data in the ENM datastore is corrupted or data loss has occurred.

In the above cases, a full restore from backup must be performed.

If a full restore is needed, refer to *Small Integrated ENM Full Restore from External Device* in Small Integrated ENM Backup and Restore System Administrator Guide.

Prerequisites

— HPE iLO access.

— Site Engineering Document (SED) is available.

— A valid recent backup is available if a restore is needed.

— Section *Post-installation Steps* in Small Integrated ENM Transport Only Installation Instructions has been completed successfully.

Note: Control stack VMs are a defined group of virtual machines:

- `VMS`
- `<vcenter_object_prefix>_VCSA`
- `management-server`
- `Openstack-Compute-0`
- `Openstack-ControlPlane-0`
- `<deployment_id>-vnflaf-services-0`
- `<deployment_id>-vnflaf-db-0`

Steps

1. Perform a platform health check before shutdown (see Perform Platform Health Check on page 13).
2. Log on to VMS as `root` with `<vms_root_password>`.
3. Disable scheduled backup policy (follow the section *Deactivate Small Integrated ENM Backup Policy to External Device or Customer-Provided NFS Share* in Small Integrated ENM Backup and Restore System Administrator Guide).
4. Stop the DDC service.

   ```
   # systemctl stop ddc.service
   ```

5. Stop the autofs service.

   ```
   # systemctl stop autofs.service
   ```

6. Run the **ENM Management** workflow to shut down ENM by following the section Shut down ENM on page 79.
7. Shut down VIO.
   a. From VMS, log on to OMS as `viouser`.

      ```
      [root@vms]# ssh viouser@oms
      [viouser@oms]#
      ```

   b. Stop the OpenStack deployment.

      ```
      [viouser@oms]# sudo viocli deployment stop
      ```

8. Shut down the control stack VMs and put the ESXi host in maintenance mode.

    a. Log on to ESXi host through SSH using `<esxi_host1_ip_vio_mgt>` as `root` with `<esxi_host1_mgt_password>`.

    Note: Do not log on to ESXi host from VMS as it has shut down during this stage and the connection is terminated.

    b. Power off the control stack VMs.

    ```
    [root@esxi:~] vim-cmd hostsvc/autostartmanager/autostop
    ```

    Note: The control stack a may take a short while to stop.

    c. List the running VMs and if the DVMS VM is listed, take note of its `World ID`.

    ```
    [root@esxi:~] esxcli vm process list
    ```

    Example

    ```
    [root@esxi:~] esxcli vm process list
    DVMS
    World ID: 75389
    Process ID: 0
    VMX Cartel ID: 75388
    UUID: 42 2f 47 26 e3 e7 09 c3-97 63 1c c1 4c 13 49 b7
    Display Name: DVMS
    Config File: /vmfs/volumes/5d9f82fc-a29ba280-f9b8-48df3714144c/DVM
    →
    S/DVMS.vmx
    ```

    d. If any control stack VMs are still running, follow the procedure in the section Autostop Not Powering Off All VMs on Small Integrated ENM Transport Only on page 141

    e. If DVMS VM is still running, gracefully shut it down using its World ID noted down in the previous step.

    ```
    [root@esxi:~] esxcli vm process kill -w <DVMS World ID> -t soft
    ```

    Note: DVMS VM may take a short while to stop.

    f. Ensure no VMs are running before the ESXi host enters the maintenance mode.

    ```
    [root@esxi:~] esxcli vm process list
    ```

    g. Request that the ESXi host enters maintenance mode.

    ```
    [root@esxi:~] esxcli system maintenanceMode set --enable true
    ```

h. Check if the system has entered maintenance mode after a short interval.

```
[root@esxi:~] esxcli system maintenanceMode get
Enabled
```

9. Power off the ESXi host through iLO console.

   a. Press **<F12>** in the iLO console window.

   b. Log in as `<esxi_host1_ilo_user>` with `<esxi_host1_mgt_password>` when prompted.

   c. Press **<F2>** to shut down the host.

Results

Small Integrated ENM Transport Only deployment is successfully shut down.

11.2 Startup Procedure for Small Integrated ENM Transport Only

> This section provides steps to start a Small Integrated ENM Transport Only deployment after it was shut down using the shutdown procedure.

Stop!

This procedure along with Shutdown Procedure for Small Integrated ENM Transport Only on page 66 cannot be used in the following scenarios:

— Hardware replacement of some or all rack components has taken place.

— The data in the ENM datastore is corrupted or data loss has occurred.

Note: In the above cases, a full restore from backup must be performed.

If a full restore is necessary, refer to the section *Small Integrated ENM Full Restore from External Device* in Small Integrated ENM Backup and Restore System Administrator Guide.

Prerequisites

— HPE iLO access.

— Site Engineering Document (SED) is available.

— A valid recent backup is available in case a restore is needed.

— This procedure can be used only when the system was previously successfully brought down with Shutdown Procedure for Small Integrated ENM Transport Only on page 66.

— Section *Post-installation Steps* in *Small Integrated ENM Transport Only Installation Instructions(1/1531-CNA 403 3456)* has been completed successfully.

Note: Control stack VMs is a defined group of virtual machines:

— VMS

— <vcenter_object_prefix>_VCSA

— management-server

— Openstack-Compute-0

— Openstack-ControlPlane-0

— <deployment_id>-vnflaf-services-0

— <deployment_id>-vnflaf-db-0

Steps

1. Power on the ESXi host through iLO console.

   Note:
   — Wait until the ESXi console indicates ESXi is loaded correctly. The ESXi finishes loading when the DCUI splash screen is available. It can take some time.

   — If ESXi fails to load, then a full restore from backup is required, refer to the section *Small Integrated ENM Full Restore from External Device* in Small Integrated ENM Backup and Restore System Administrator Guide.

2. Bring the ESXi host out of maintenance mode and start the control stack up.

   a. Log on into ESXi host through <esxi_host1_ip_vio_mgt> as root with <esxi_host1_mgt_password>.

   b. Request ESXi host to exit maintenance mode.

   ```
   # esxcli system maintenanceMode set --enable false
   ```

   c. Power on the control stack VMs.

   ```
   # vim-cmd hostsvc/autostartmanager/autostart
   ```

Do!

Wait for 20 minutes to allow the control stack VMs to come up fully running.

3. Check that the control stack VMs are back and powered on.

    a. Open a web browser to the ESXi host client `https://<esxi_host1_ip_vio_mgt>` and log in as user `root` with password `<esxi_host1_mgt_password>`.

    b. Look for the control stack VMs, right-click, select **Console** and select **Open browser console**.

    c. Verify the state of each control stack VM except for VNF-LAF VMs as described in the table *Control Stack VMs Status in Console*.

Table 3 Control Stack VMs Status in Console

Control Stack VM	Status in Console
`VMS`	VMS login prompt
`<vcenter_object_prefix>_VCSA`	DCUI splash screen or operating system login prompt
`management-server`	DCUI splash screen or operating system login prompt
`Openstack-Compute-0`	`compute01` login prompt
`Openstack-ControlPlane-0`	`loadbalancer01` login prompt

4. Power on the DVMS VM (in the **VMs** tab, right-click **DVMS** and select **Power > Power On**).

5. Ensure the `vapi` service is running on the OMS.

    a. SSH to the VMS on the management network IP `<vms_ip_vio_mgt>` as `root` with password `<vms_root_password>` from the ENM SED.

    b. SSH to the OMS from the VMS as user `viouser`; switch to `root` user; and restart the `vapi` service.

```
[root@vms] # ssh viouser@oms
[viouser@oms] # sudo -i
[root@oms] # systemctl start vapi
```

6. Perform a VIO platform health check (see section Perform Platform Health Check on page 13).

7. Start ENM Management workflow (see section Start ENM on page 81).

8. Perform ENM System health check (see section ENM Health Check on page 22).

9. Enable scheduled backup policy (see section *Reactivate Small Integrated ENM Backup Policy to External Device or Customer-Provided NFS Share* in Small Integrated ENM Backup and Restore System Administrator Guide.

Results

Small Integrated ENM Transport Only deployment is fully operational.

11.3 Recover from an Unplanned Outage of Small Integrated ENM Transport Only

> This section provides steps to recover from an unplanned outage of a Small Integrated ENM Transport Only deployment.

Stop!

This recovery procedure cannot be used in the following scenarios:

— Hardware replacement of some or all rack components has taken place.

— The data in the ENM datastore is corrupted or data loss has occurred.

In these scenarios, a full restore from backup must be performed.

If a full restore is necessary, refer to the section *Small Integrated ENM Full Restore from External Device* in *ENM on Cloud Upgrade Instructions (2/153 72-AOM 901 151).*

Prerequisites

— HPE iLO access.

— Site Engineering Document (SED) is available.

— No network or hardware faults exist.

— Section *Post-installation Steps* in *Small Integrated ENM Transport Only Installation Instructions (1/1531-CNA 403 3456)* has been completed successfully.

Stop!

If the below procedure is unsuccessful, then a full restore from backup is required, refer to the section *Small Integrated ENM Full Restore from External Device* in Small Integrated ENM Backup and Restore System Administrator Guide.

Steps

1. Recover the ESXi host.

    a. Use the iLO console to power on the server if it is powered off (if the server is already powered on then reset the power through the iLO console).

    Note:
    — Wait until the ESXi console indicates ESXi is loaded.

    — If ESXi fails to load then a full restore from backup is required, see section *Small Integrated ENM Full Restore from External Device* in Small Integrated ENM Backup and Restore System Administrator Guide.

    — If the ESXi GUI has a connection timeout at any stage, or if the `vim-cmd` hangs while getting information about running VMs, then restart the `hostd` service with the following command.

    ```
    # /etc/init.d/hostd restart
    ```

    b. Open a web browser to the ESXi host client `https://<esxi_host1_ip_vio_mgt>` and log on as user `root` with password `<esxi_host1_mgt_password>`.

    c. Click **Storage** in the **Navigator** pane and check that `datastore1` is available.

    d. Check if the `datastore1` is healthy by selecting the `datastore1` available under **Storage** in **Navigator** pane and click the **Datastore browser**.

    Note: A list of files and folders should be visible in the browser.

    e. Check if the ESXi host is in maintenance mode.

    Note: If the ESXi host is in maintenance mode then exit maintenance mode by right-clicking the **Host** in the **Navigator** pane and click **Exit maintenance mode**.

    f. Connect through SSH on to the ESXi host `<esxi_host1_vio_mgt_hostname>` as user `root` with password

<esxi_host1_mgt_password> from the SED and run the following command to auto-start the following control stack VMs:

- VMS
- <vcenter_object_prefix>_VCSA
- management-server
- Openstack-Compute-0
- Openstack-ControlPlane-0
- <deployment_id>-vnflaf-services-0
- <deployment_id>-vnflaf-db-0

```
# vim-cmd hostsvc/autostartmanager/autostart
```

2. Check control stack VMs are back and powered on.

Do!

Wait for 20 minutes until the control stack VMs are fully running.

a. Click **Virtual Machines** in the previously opened ESXi GUI through the **Navigator** pane.

b. Look for the control stack VMs, right click, select Console and select Open browser console.

c. Verify the state of each control stack VMs as described in the table *Control stack VMs status in console*.

Table 4 Control Stack VMs Status in Console

Control Stack VM	Status in Console
VMS	VMS login prompt
<vcenter_object_prefix>_VCSA	DCUI splash screen or operating system login prompt
management-server	DCUI splash screen or operating system login prompt
Openstack-Compute-0	compute01 login prompt
Openstack-ControlPlane-0	loadbalancer01 login prompt
<deployment_id>-vnflaf-services-0	<deployment_id>_vnflaf-services login prompt
<deployment_id>-vnflaf-db-0	<deployment_id>_vnflaf-db login prompt

Note:

— If any of the control stack VMs failed to boot up because of file system errors, similar to the error shown, follow the section Resolve File System Errors on page 134.

```
/dev/sda6 contains a file system with errors, check fo →
rced.
/dev/sda6: Inodes that were part of a corrupted orphan →
 linked list found.
/dev/sda6: UNEXPECTED INCONSISTENCY; RUN fsck MANUALLY →
.
(i.e., without -a or -p options)
fsck exited with status code 4
The root filesystem on /dev/sda6 requires a manual fsc →
k
```

- Where `/dev/sda6` in the example is the file system of the VM that contains the error.
- The file systems which contain the error can be identified from the error shown in the console.

— Wait for 20 minutes to allow the control stack VMs to be running.

3. Check that the external backup device is recognized and mounted on the VMS.

   a. Check that the external backup device is recognized and mounted on the VMS.

   b. Connect over SSH to the VMS management network IP `<vms_ip_vio_mgt>` and log in as user `root` with password `<vms_root_password>`.

   c. Run the following command and find the external backup device which is mounted to the VMS.

   ```
   # mount | grep BACKUP
   ```

   Example

   ```
   [root@vms ~]# mount | grep BACKUP
   /dev/sdd on /BACKUP type ext4 (rw,relatime,seclabel,data=ordered)
   ```

   d. Run the following command to activate the external backup device that is found after running the preceding command:

   ```
   # mount -a
   ```

4. Disable the scheduled backup policy by following section *Deactivate Small Integrated ENM Backup Policy to External Device or Customer-Provided NFS Share* in Small Integrated ENM Backup and Restore System Administrator Guide).

5. Ensure the `vapi` service is running on the OMS.

    a. Connect over SSH to the VMS on the management network IP `<vms_ip_vio_mgt>` as `root` with password `<vms_root_password>` from the ENM SED.

    b. Connect over SSH to the OMS from the VMS as user `viouser`; switch to `root` user; and restart the `vapi` service.

    ```
    [root@vms] # ssh viouser@oms
    [viouser@oms] # sudo -i
    [root@oms] # systemctl start vapi
    ```

6. Perform VIO Platform Health Check by following section Perform Platform Health Check on page 13.

7. Run ENM Management workflow to recover ENM by following section Recover ENM on page 85.

8. Perform an ENM System Health Check by following section ENM Health Check on page 22.

9. Enable the previously disabled schedule backup policy by following the section *Reactivate Small Integrated ENM Backup Policy to External Device or Customer-Provided NFS Share* in Small Integrated ENM Backup and Restore System Administrator Guide.

Results

Small Integrated ENM Transport Only deployment is fully recovered after an unplanned outage.

12 Manage ENM

> The ENM on Cloud Management workflows allow stop, start, and recovery of ENM on Small Integrated ENM Multi-Technology and Transport deployments.

Transport deployments use the workflows for planned maintenance and recovery of ENM after an unplanned loss of the ESXi host.

Multi-Technology deployments use the workflows for planned maintenance on two or more ESXi hosts in the cluster and to recover ENM after unplanned loss of two or more hosts.

Note: Multi-Technology does not need to shut down ENM for planned maintenance or ENM recovery after loss of a single host as ENM can run on the remaining two ESXi hosts.

Prerequisites

- OpenStack RC file available
- `sed.json` is in the `/vnflcm-ext/enm/` directory on the VNF-LAF Services VM.
- VMS is installed.
- VIO is deployed.
- VNF-LCM is deployed.
- `ERICenmdeploymentworkflows_CXP9034151` RPM installed on VNF-LCM.
- `ERICenmcloudmgmtworkflows_CXP9036442` RPM installed on VNF-LCM.
- `ERICenmcloudmgmtutils_CXP9036444` RPM installed on VMS.

Troubleshooting

For information on known issues and troubleshooting, see Troubleshoot Problems with ENM Management on page 129.

12.1 Execute Manage ENM on the VMS

The following section goes through how to manage ENM on the VMS.

12.1.1 Manage ENM Script Usage

This section provides a general overview of `manage_enm` script usage.

The `manage_enm` script provides an automated solution for running the ENM on Cloud Management workflows.

Steps

1. Log on to the VMS as `root`.
2. Check that the script usage.

```
[root@vms ~]# manage_enm --help

INFO: Logging to /root/manage_enm.log
usage: manage_enm.py [-h] [--rcfile RCFILE] [--lcm HOST] [--lcm-name LCM_NAM →
E]
                     [--lcm-db-name LCM_DB_NAME]
                     {stop,start,recover} ...

Start or stop ENM. By default, the shutdown procedure is graceful.

positional arguments:
  {stop,start,recover}
    stop                Shut down ENM
    start               Start ENM
    recover             Recover ENM by stopping/starting

optional arguments:
  -h, --help            show this help message and exit
  --rcfile RCFILE       path to the OpenStack RC file
  --lcm HOST            hostname or IP of VNF-LCM workflow service
  --lcm-name LCM_NAME   Custom value for VNF-LCM services server as defined  →
in VNF-LCM SED.
  --lcm-db-name LCM_DB_NAME
                        Custom value for VNF-LCM DB server as defined in VNF →
-LCM SED.

examples:
manage_enm stop
manage_enm stop --hard --reason "Firmware updates" --skiplcm --yes
manage_enm stop --lcm-name lcmName --lcm-db-name lcmDBName

manage_enm start --rcfile /tmp/file.rc --lcm 10.10.22.22
manage_enm recover --reason "Unexpected loss of host"
```

3. When starting the script, actions can be combined with options.
   - If the OpenStack RC file is not sourced, then it can be specified as an optional argument.

```
[root@vms ~]# manage_enm stop --rcfile /vol1/senm/etc/vio123_project.rc
```

— The internal IP address of VNF-LCM LAF Services VM can be specified as an optional argument.

Note: By default the script uses the hostname `vnflaf-services`.

```
[root@vms ~]# manage_enm start --lcm 10.10.10.123
```

— The script uses `vnflaf-services` and `vnflaf-db` names to start, stop, and recover the VNF-LCM VMs by default. If the properties `<Services_vm_HostName>` and `<DB_vm_HostName>` in the VNF-LCM SED are customized, provide the customized values using `--lcm-name` and `--lcm-db-name` arguments.

```
[root@vms ~]# manage_enm stop --lcm-name custom-vmserv-name --lcm-db-na →
me custom-vmdb-name
```

4. The `manage_enm` script creates a log file called `manage_enm.log` in the users home directory.

12.1.2 Shut down ENM

> This section describes how to use the `manage_enm` script to shut down ENM before planned maintenance.

The script runs the **Manage ENM - Shutdown** workflow which deletes all ENM stacks except the VNF-LCM and the `network_security_group` stacks.

The `manage_enm` script performs either a graceful or a hard shutdown of ENM.

Note:
— The graceful option shuts down the VMs before deleting stacks while hard shutdown deletes stacks without shutting down VMs beforehand.

— Volumes are retained in both cases.

— With every shutdown an FM alarm is generated with a default reason for shutdown: `Shutdown ENM for planned maintenance`.

— The approximate time to complete the shutdown workflow with the graceful option is 10 minutes and, for a hard shutdown, 3 minutes.

Steps

1. Source the OpenStack RC file for the ENM deployment.

```
[root@vms ~]# . /vol1/senm/etc/vio_123_project.rc
```

2. Execute the shutdown with the default option (graceful) or with the hard shutdown option.

   — Graceful shutdown (Default).

   ```
   [root@vms ~]# manage_enm stop
   ```

   — Hard shutdown

   ```
   [root@vms ~]# manage_enm stop --hard
   ```

   Example

   ```
   [root@vms ~]# manage_enm stop
   INFO: Logging to /root/manage_enm.log

   This procedure will shut down ENM, causing unavailability of services. Type →
   "YeS" when you are ready to proceed:  YeS
   INFO: Authenticating with OpenStack
   INFO: Successfully authenticated with OpenStack
   INFO: Checking if essential OpenStack services are up
   INFO: Checking if VNF-LCM workflows are available at "vnflaf-services"
   INFO: Stopping ENM with graceful option. Reason: manage_enm - Shutdown ENM f →
   or planned maintenance
   INFO: Starting ShutdownENM__top workflow
   INFO: Monitoring workflow instance ShutdownENM_20181121_135117
   ......
   INFO: Workflow no longer running
   INFO: Workflow instance completed successfully
   ```

   Note:
   — The user prompt can be omitted by adding `--yes` option to the `manage_enm stop` command.

   — If the `manage_enm` script fails with a connection error, refer to the section Manage ENM Script Fails with 'ConnectionError' on page 131 before continuing with the VIO deployment shutdown.

   — If the script fails, do not proceed to the next step, collect the `manage_enm.log` file located under the root directory and contact Ericsson support.

3. Log onto the VMS using the *<vms_ip_vio_mgt>* IP address as user `root` with password *<vms_root_password>*, where *<vms_ip_vio_mgt>* and *<vms_root_password>* are parameters in the ENM SED.

4. Log onto the OMS VM.

   ```
   # ssh viouser@oms
   ```

5. Log onto the `loadbalancer` VM.

   ```
   $ ssh loadbalancer01
   ```

6. Monitor the system log by running the following command until VIO releases all IP addresses. This may take up to 5 minutes. The command queries the

log every second. When VIO completes releasing of IP addresses, the new DHCPRELEASE entries stop appearing in the system log. The system log can be empty, if VIO has released all the IP addresses.

```
$ watch -n1 grep -i DHCPRELEASE /var/log/syslog
```

Example

System Log with DHCPRELEASE Entries

```
$ watch -n1 grep -i DHCPRELEASE /var/log/syslog
Every 1.0s: grep -i DHCPRELEASE /var/log/syslog                                 →
                                          Thu Nov 26 14:33:06 2020

Nov 26 14:12:46 loadbalancer01 dnsmasq-dhcp[22343]: DHCPRELEASE(tap05d66c8e-   →
bb) 10.10.2.19 fa:16:3e:50:89:ba
Nov 26 14:12:47 loadbalancer01 dnsmasq-dhcp[22343]: DHCPRELEASE(tap05d66c8e-   →
bb) 10.10.2.18 fa:16:3e:6e:82:dc
Nov 26 14:12:48 loadbalancer01 dnsmasq-dhcp[22343]: DHCPRELEASE(tap05d66c8e-   →
bb) 10.10.2.37 fa:16:3e:e4:5b:60
Nov 26 14:12:49 loadbalancer01 dnsmasq-dhcp[22343]: DHCPRELEASE(tap05d66c8e-   →
bb) 10.10.2.127 fa:16:3e:02:27:35
Nov 26 14:12:50 loadbalancer01 dnsmasq-dhcp[22343]: DHCPRELEASE(tap05d66c8e-   →
bb) 10.10.2.147 fa:16:3e:53:3b:45
Nov 26 14:12:51 loadbalancer01 dnsmasq-dhcp[22343]: DHCPRELEASE(tap05d66c8e-   →
bb) 10.10.2.169 fa:16:3e:e7:2e:80
Nov 26 14:12:53 loadbalancer01 dnsmasq-dhcp[22343]: DHCPRELEASE(tap05d66c8e-   →
bb) 10.10.2.137 fa:16:3e:16:6d:56
Nov 26 14:12:54 loadbalancer01 dnsmasq-dhcp[22343]: DHCPRELEASE(tap05d66c8e-   →
bb) 10.10.2.10 fa:16:3e:50:dd:49
Nov 26 14:12:55 loadbalancer01 dnsmasq-dhcp[22343]: DHCPRELEASE(tap05d66c8e-   →
bb) 10.10.2.66 fa:16:3e:64:d9:93
Nov 26 14:12:56 loadbalancer01 dnsmasq-dhcp[22343]: DHCPRELEASE(tap05d66c8e-   →
bb) 10.10.2.38 fa:16:3e:1b:31:ae
Nov 26 14:12:57 loadbalancer01 dnsmasq-dhcp[22343]: DHCPRELEASE(tap05d66c8e-   →
bb) 10.10.2.168 fa:16:3e:bc:c0:d6
Nov 26 14:12:58 loadbalancer01 dnsmasq-dhcp[22343]: DHCPRELEASE(tap05d66c8e-   →
bb) 10.10.2.166 fa:16:3e:a5:88:fd
```

Example

System Log with no DHCPRELEASE Entries

```
$ watch -n1 grep -i DHCPRELEASE /var/log/syslog
Every 1.0s: grep -i DHCPRELEASE /var/log/syslog                                 →
                                          Thu Nov 26 15:18:50 2020
```

Results

ENM is successfully shut down.

12.1.3 Start ENM

> This section describes how to use the manage_enm script to start ENM after planned maintenance.

The script runs the **Manage ENM - Start** workflow which recreates the ENM stacks using the existing volumes.

Note:

— The workflow fails if there are any existing ENM stacks or if there are any volumes which are not in available state. Only the VNF-LCM and the `network_security_group` stacks should be present.

— The approximate time to complete the start procedure is 30–40 minutes.

Prerequisites

ENM has been shut down by following Shut down ENM on page 79.

Steps

1. Log on to the VMS as `root`.
2. Source the OpenStack RC for the ENM deployment.

   Example

   ```
   [root@vms ~]# . /vol1/senm/etc/vio_123_project.rc
   ```

3. Check if the volumes are in `available` state. If any volumes are in `reserved` state, then reset it to `available` state.

   a. Check if all ENM volumes have state `available` except VNF-LCM.

   ```
   [root@vms]# openstack volume list
   ```

   Example

   ```
   [root@vms]# openstack volume list
   +--------------------------------------+--------------------------------
   -+-----------+------+-----------------------------------+
   | ID                                   | Display Name                   
    | Status    | Size | Attached to                       |
   +--------------------------------------+--------------------------------
   -+-----------+------+-----------------------------------+
   | f43669f0-70b4-40a6-b33b-             | vio-5604-nfssmrs_volume-0      
    | available |  100 |                                   |
   | 289d1745e32c                         |                                
   |              |      |                                   |
   | bebef49f-                            | vio-5604-nfspm_volume-0        
    | available |  350 |                                   |
   | e708df26-0c3c-4adf-bec4-             |                                
   |              |      |                                   |
   | ec7531cda199                         | vio-5604-ops_volume-1          
    |  reserved |    1 |                                   |
   
   | ........                             |.......                         
    |.......    |..... |.......                            |
   | ........                             |.......                         
    |.......    |..... |.......                            |
   | ........                             |.......                         
    |.......    |..... |.......                            |
   | ........                             |.......                         
    |.......    |..... |.......                            |
   | 8b3b5363-4552-4371-afd3-a78b8b5      |                                
   |              |      |                                   |
   | 0db52                                | vio-5604-vnflcm-volume-0       
    |   in-use  |  120 | Attached to vio-5604-vnflaf       |
   ```

```
+--------------------------------+--------------------------------
-+-----------+------+---------------------------------+
```

b. Check for the volumes in state `reserved`.

```
[root@vms]# openstack volume list | grep -i reserved
```

Example

```
[root@vms]# openstack volume list | grep -i reserved
| e708df26-0c3c-4adf-bec4-ec7531cda199 | vio-5604-ops-volume-1 | re
served | 1 | |
```

c. If there are any reserved volumes found in the previous step, then reset the reserved volume state to **`detached and available`**.

```
[root@vms]# cinder reset-state --attach-status detached --state ava
ilable <Volume ID from above output>
```

Example

```
[root@vms]# cinder reset-state --attach-status detached --state ava
ilable e708df26-0c3c-4adf-bec4-ec7531cda199
```

Note: Repeat the steps for all the volumes found in reserved state.

4. Start ENM using the `manage_enm` script (the script runs the **Manage ENM - Start** workflow to recreate the stacks for ENM).

```
[root@vms ~]# manage_enm start
```

Example

```
[root@vms ~]# manage_enm start
INFO: Logging to /root/manage_enm.log
INFO: Authenticating with OpenStack
INFO: Successfully authenticated with OpenStack
INFO: Checking if essential OpenStack services are up
INFO: Checking if VNF-LCM application is available at "vnflaf-services".
INFO: Checking if workflow StartENM__top is installed
INFO: Starting workflow with ID: enmcloudmgmtworkflows.--.1.0.5.--.StartENM_
_top
INFO: Workflow started with business key StartENM_20180716_115104
INFO: Beginning to monitor workflow execution
INFO: Workflow no longer active
INFO: Workflow StartENM__top has completed successfully
[root@vms ~]#
```

Note:
- HA workflows may be observed for one or more instances during the ENM startup phase of the recovery or startup. Alarms relating to High Host CPU Usage may also be observed in vCenter UI.This is normal for deployments managing many network elements.
- If the HA workflows and vCenter alarms are ongoing several hours after the ENM recovery or startup is complete, contact Ericsson support.

5. Open ENM Launcher in your browser, using the value of `<httpd_fqdn>` specified in the SED and log in. The list of applications should be available.

6. Check all stacks are in `CREATE_COMPLETE` state.

```
[root@vms ~]# openstack stack list -f value -c "Stack Status" | grep -v CREA →
TE_COMPLETE
```

Note: This command returns no output if all stacks are in `CREATE_COMPLETE`.

7. Check that all VMs in OpenStack are in Consul.

   a. Find the number of VMs in Consul.

   Note: Use the value of `<key_name>` defined in the SED.

```
[root@vms ~]# ssh -i /vol1/senm/etc/<key_name>.pem cloud-user@vnfla →
f-services 'consul members | grep -v Node | wc -l'
```

   Example

```
[root@vms ~]# ssh -i /vol1/senm/etc/key_pair_vio_5592.pem cloud-use →
r@vnflaf-services 'consul members | grep -v Node | wc -l'
164
```

   b. Find the number of VMs in OpenStack.

```
[root@vms ~]# openstack server list -f value -c Name | grep -iv vnf →
laf-db | wc -l
```

   Note: The `vnflaf-db` VM is excluded as the Consul agent does not run on this VM.

   Example

```
[root@vms ~]# openstack server list -f value -c Name | grep -iv vnf →
laf-db | wc -l
164
```

   Note: The number of VMs in Consul should match the number of VMs in OpenStack (excluding the `vnflaf-db` VM).

Results

ENM is started successfully.

12.1.4 Recover ENM

> This section describes how to use the `manage_enm` script to recover ENM after an unplanned loss of the ESXi host (Transport) or of two or more ESXi hosts (Multi-Technology).

The script recovers the VNF-LCM service if it is not operational. It then triggers the shutdown and start workflows. The **Manage ENM - Shutdown** workflow is executed with the default (graceful) option which shuts down VMs before deleting stacks. To reduce the load on the deployment, the `manage_enm` script waits five minutes before triggering the **Manage ENM - Start** workflow. The approximate time to recover ENM is 45 minutes.

Steps

1. Log on to the VMS as `root`.
2. Source the OpenStack RC for the ENM deployment.

   Example

   ```
   [root@vms ~]# source /vol1/senm/etc/vio_123_project.rc
   ```

3. Execute the `manage_enm` script with the `recover` option.

   Example

   ```
   [root@vms ~]# manage_enm recover --reason "Recover VIO after loss of power"
   INFO: Logging to /root/manage_enm.log
   INFO: Recover ENM started
   INFO: Authenticating with OpenStack
   INFO: Checking if essential OpenStack services are up
   INFO: Checking if VNF-LCM workflows are available at "vnflaf-services"
   INFO: Starting ShutdownENM__top workflow
   INFO: Monitoring workflow instance ShutdownENM_20181011_101425
   ............
   INFO: Workflow no longer running
   INFO: Workflow instance completed successfully
   INFO: Waiting 300 seconds before starting ENM
   INFO: Authenticating with OpenStack
   INFO: Checking if essential OpenStack services are up
   INFO: Checking if VNF-LCM application is available at "vnflaf-services"
   INFO: Starting ENM
   INFO: Starting StartENM__top workflow
   INFO: Monitoring workflow instance StartENM_20181011_102557
   ..............................................................
   INFO: Workflow no longer running
   INFO: Workflow instance completed successfully
   ```

Note:
- Providing a reason for recovery is optional. If not provided, the default reason `manage_enm - Recover ENM after unexpected downtime` is passed to **Manage ENM - Shutdown** workflow.
- HA workflows may be observed for one or more instances during the ENM startup phase of the recovery or startup. Alarms relating to `High Host CPU Usage` may also be observed in vCenter UI.
- The above is normal for deployments managing many network elements. If the HA workflows and vCenter alarms are ongoing a number of hours after the ENM recovery or startup is complete, contact Ericsson support.

Results

ENM is successfully recovered.

12.2 Manage ENM with Workflows on VNF-LCM

> This section describes the procedures required to shut down or start ENM using workflows on VNF-LCM.

Note: ENM is managed using the workflows, and VNF-LCM remains powered on.

12.2.1 Shut Down ENM on Cloud using VNF-LCM

> This section describes how to shut down ENM on Cloud using VNF-LCM.

This is achieved by running the **Manage ENM - Shutdown** workflow which performs either a graceful or hard shutdown of ENM on Cloud.

Note:
- The graceful option shuts down VMs before deleting stacks, while the hard shutdown deletes stacks without shutting down VMs beforehand. Volumes are retained in both cases.
- Approximate time to complete shutdown with the graceful option is 10 minutes and, for the hard option, 3 minutes.

Prerequisites

- ENM on Cloud deployed.
- VNF-LCM deployed.

— VNF-LCM SED is available.

— `ERICenmcloudmgmtworkflows_CXP9036442` rpm is installed on VNF-LCM.

— `ERICenmdeploymentworkflows_CXP9034151` rpm installed on VNF-LCM.

Steps

1. Open the VNF-LCM UI using the following URL.

   ```
   http://<external_ip_for_services_vm>/index.html#workflows
   ```

   Note: Replace the value for `<external_ip_for_services_vm>` with the value corresponding to either the `<external_ipv4_for_services_vm>` or `<external_ipv6_for_services_vm>` parameter in the VNF-LCM SED.

2. Select the **Manage ENM - Shutdown** workflow, then click **Start a New Instance** (alternatively, right-click **Manage ENM - Shutdown** workflow and select **Start a New Instance**).

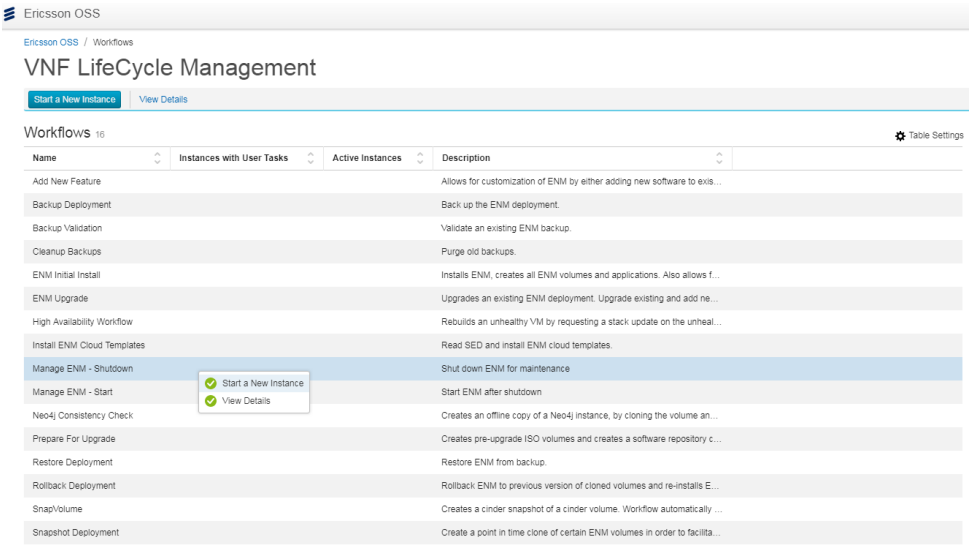

3. Select either **Graceful Shutdown** or **Hard Shutdown**, then click **Submit** to start the ENM on Cloud shutdown.

   Example

   A graceful **Manage ENM - Shutdown**.

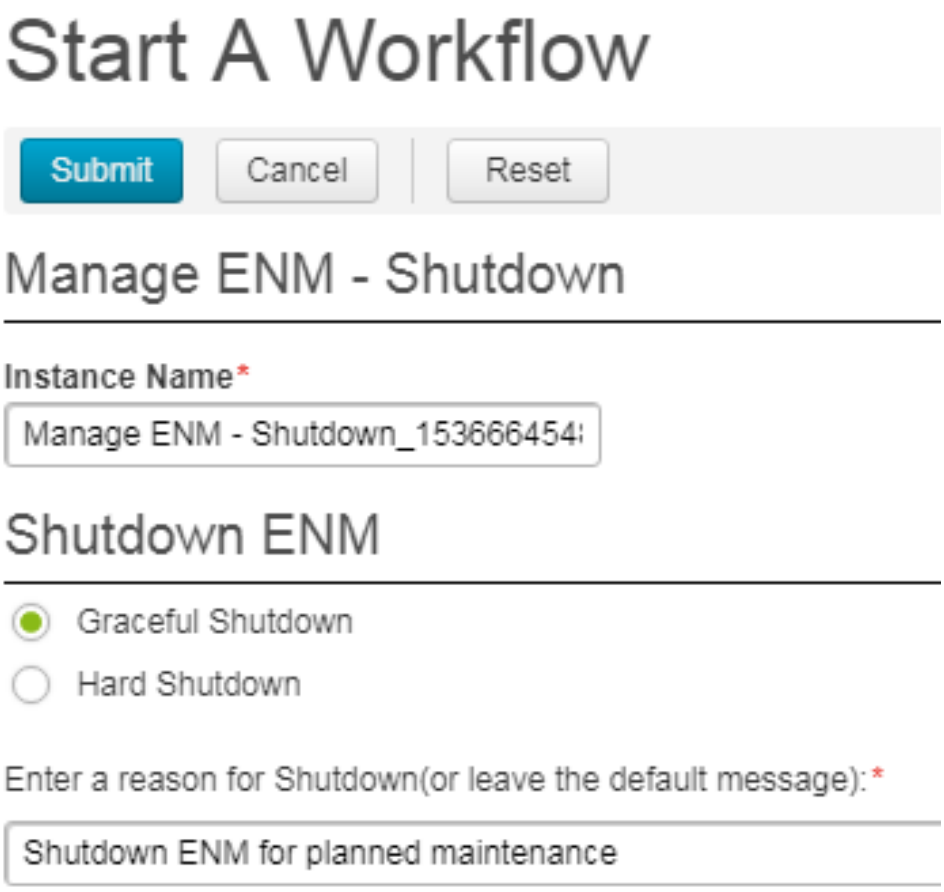

Note: The shutdown reason field can be edited to provide a more detailed reason for the shutdown. This value will be used by the FM alarm in the `Probable Cause` field.

4. Monitor the **Manage ENM - Shutdown** workflow (use the refresh option to reload and check the progress of the workflow instance).

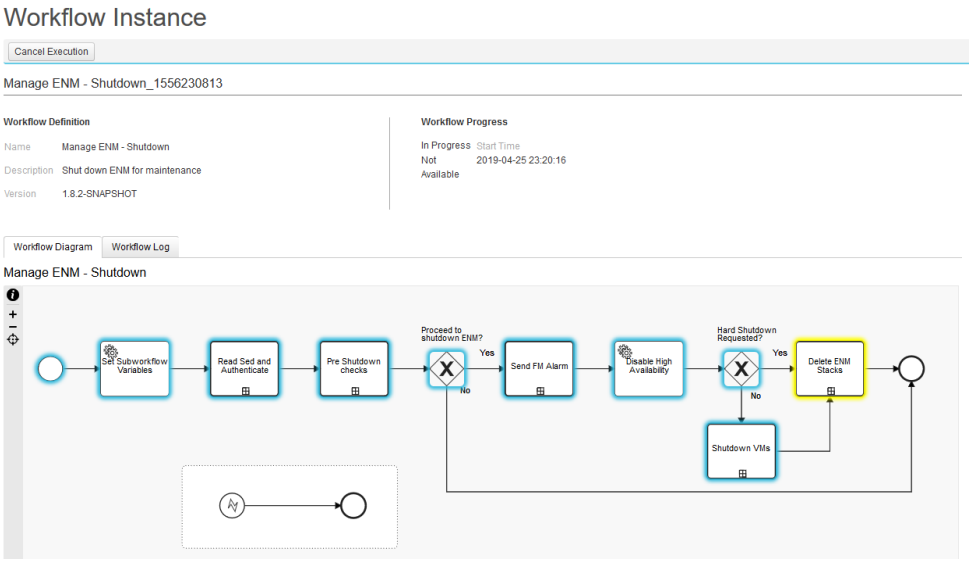

5. Wait for the completion of the **Manage ENM - Shutdown** workflow.

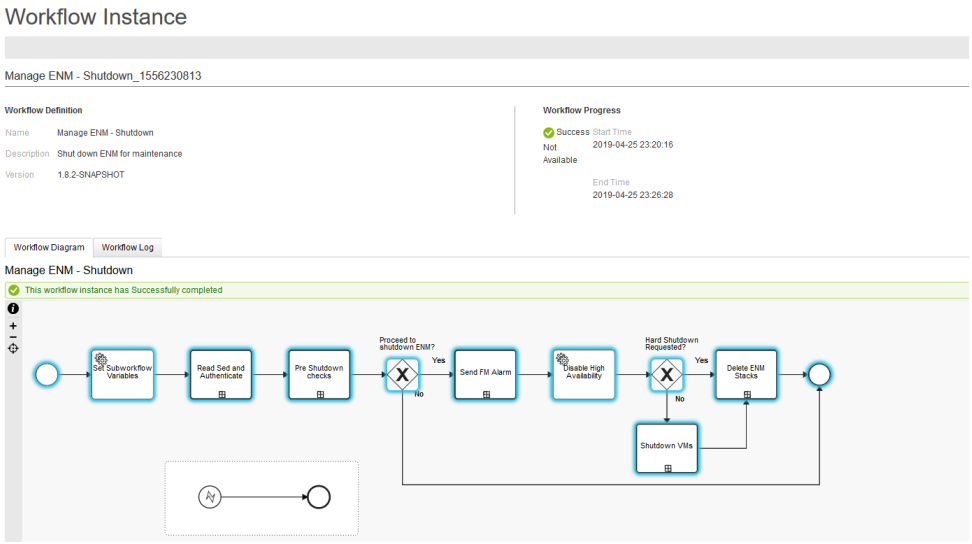

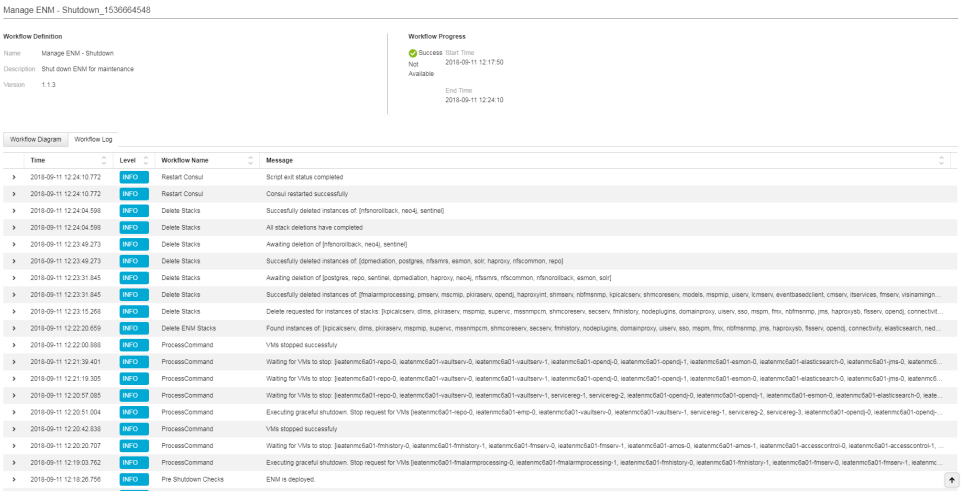

6. Gather information about the state of the system after the shutdown.

    a. Source the OpenStack RC file from the OpenStack client machine for the ENM on Cloud deployment.

    ```
    # source <OpenStack RC file name>.rc
    ```

    b. Run the following OpenStack commands to verify the deployment state.

    ```
    # openstack stack list
    # openstack server list
    # openstack volume list
    ```

    Note: There should be no ENM on Cloud stacks or VMs, only volumes. The VNF-LAF services and database VMs, stacks, and volumes should still be present.

    Example

    ```
    [root@vms ~]# openstack stack list --short

    +--------------------------------------+--------------------------- →
    -----+-----------------+
    | ID                                   | Stack Name                 →
         | Stack Status    |
    +--------------------------------------+--------------------------- →
    -----+-----------------+
    | 6c3fcfca-46e3-4dd2-a887-e0ea4ff7258b | vio567_network_security_gr →
    oup  | CREATE_COMPLETE |
    | 668dbe5f-8041-4fd7-a822-c0592ce000fc | vio567_VNFLCM              →
         | CREATE_COMPLETE |
    | 02c764fa-6008-4749-a37f-d193ae8c2c68 | vio567_vnflcm_security_gro →
    up   | CREATE_COMPLETE |
    +--------------------------------------+--------------------------- →
    -----+-----------------+

    [root@vms ~]# openstack server list -c ID -c Name -c Status
    +--------------------------------------+--------------------------- →
    -+--------+
    | ID                                   | Name                       →
     | Status |
    +--------------------------------------+--------------------------- →
    -+--------+
    ```

```
| 9dc0563e-351f-4bc3-bf2d-4497d5615454 | vio567-vnflaf-services-0
 | ACTIVE |
| ab073b0f-8535-410e-b8da-ab2179b1fe1a | vio567-vnflaf-db-0
 | ACTIVE |
+--------------------------------------+--------------------------
-+---------+

[root@vms ~]# openstack volume  list
+--------------------------------------+--------------------------
-+---------+------+---------------------------------------------+
| ID                                   | Display Name
 | Status  | Size | Attached to                                 |
+--------------------------------------+--------------------------
-+---------+------+---------------------------------------------+
/.../
| 1bd643b7-3618-4cac-b146-7286886a9e42 | vio567-rhel7_iso_volume-0
 |available|    5 |                                             |
| 098774d5-7c6c-4ead-bea3-d6b09fccb4b3 | vio567-rhel6vol-0
 |available|    5 |                                             |
| c7c64d06-6ed7-44c8-aff7-649375c2f11d | vio567_vnflcm_volume
 | in-use  |  120 | Attached to vio567-vnflaf-db-0 on /dev/sdb|
/.../
+--------------------------------------+--------------------------
-+---------+------+---------------------------------------------+
```

7. Log onto the VMS using the *<vms_ip_vio_mgt>* IP address as user `root` with password *<vms_root_password>*, where *<vms_ip_vio_mgt>* and *<vms_root_password>* are parameters in the ENM SED.

8. Log onto the OMS VM.

```
# ssh viouser@oms
```

9. Log onto the `loadbalancer` VM.

```
$ ssh loadbalancer01
```

10. Monitor the system log using the command below until VIO releases all IP addresses; the command queries the log every second. VIO has released IP addresses once new `DHCPRELEASE` entries stop appearing in the system log.

```
$ watch -n1 grep -i DHCPRELEASE /var/log/syslog
```

12.2.2 Start ENM on Cloud Using VNF-LCM

> This section describes how to start ENM on Cloud using VNF-LCM.

This is achieved by running the **Manage ENM - Start** workflow.

Approximate time to complete the start procedure is 30–40 minutes.

Prerequisites

— ENM on Cloud deployed.

— VNF-LCM deployed.

— VNF-LCM SED is available.

— `ERICenmcloudmgmtworkflows_CXP9036442` rpm is installed on VNF-LCM.

— `ERICenmdeploymentworkflows_CXP9034151` rpm is installed on VNF-LCM.

— ENM has been successfully shut down using the **Manage ENM - Shutdown** workflow.

Steps

1. Open the VNF-LCM UI using the following URL.

   ```
   http://<external_ip_for_services_vm>/index.html#workflows
   ```

   Note: Replace the value for `<external_ip_for_services_vm>` with the value corresponding for either the `<external_ipv4_for_services_vm>` or `<external_ipv6_for_services_vm>` parameter in the VNF-LCM SED.

2. Select the **Manage ENM - Start** workflow, then click **Start a New Instance** (alternatively, right-click the **Manage ENM - Start** workflow and select **Start a New Instance**).

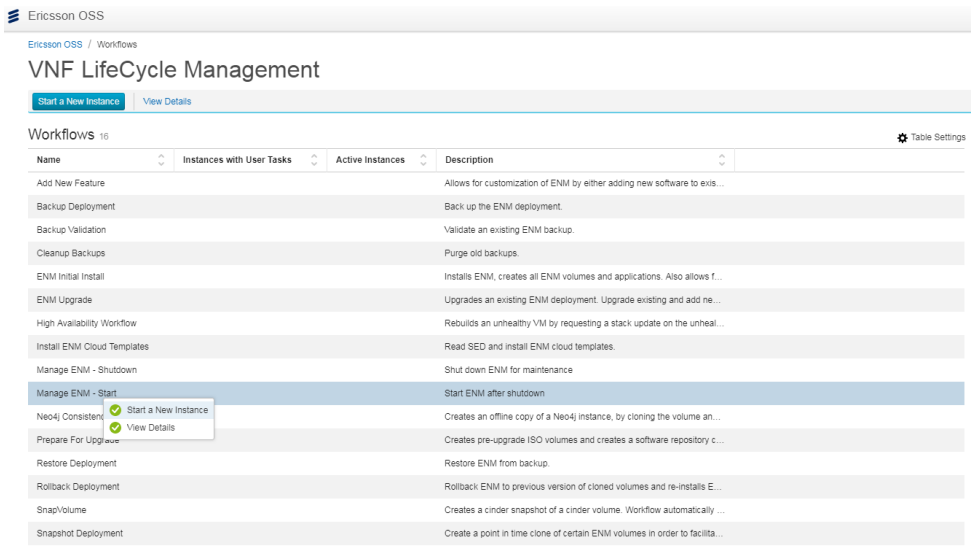

3. Click **Submit** to start the **Manage ENM - Start** workflow.

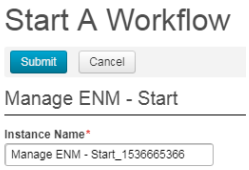

4. Monitor the **Manage ENM - Start** workflow (use the refresh option to reload and check the progress of the workflow instance).

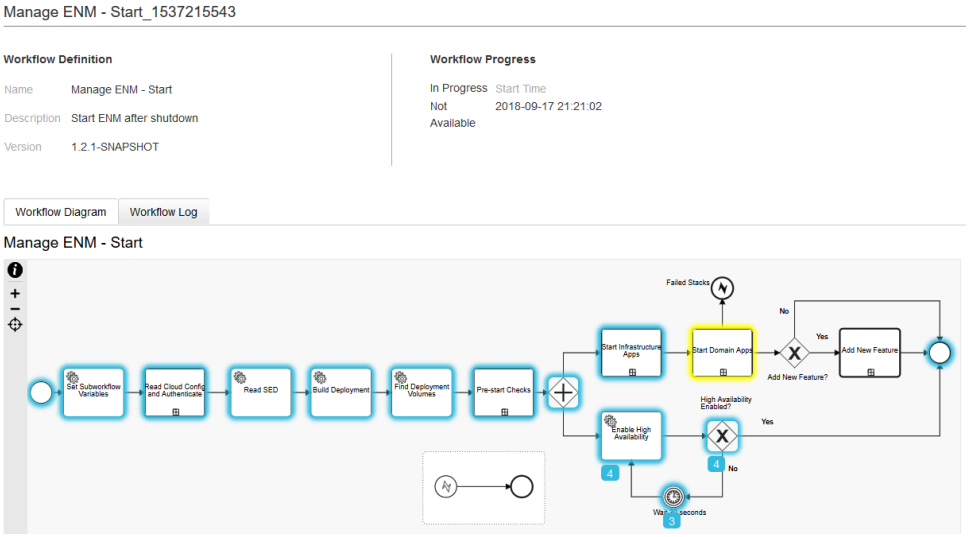

5. Wait for the completion of the **Manage ENM - Start** workflow.

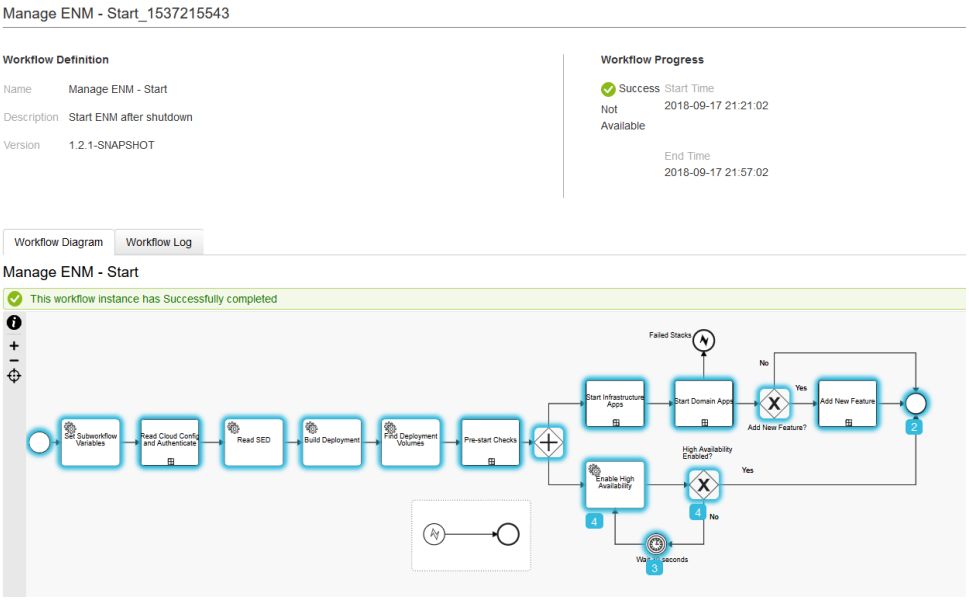

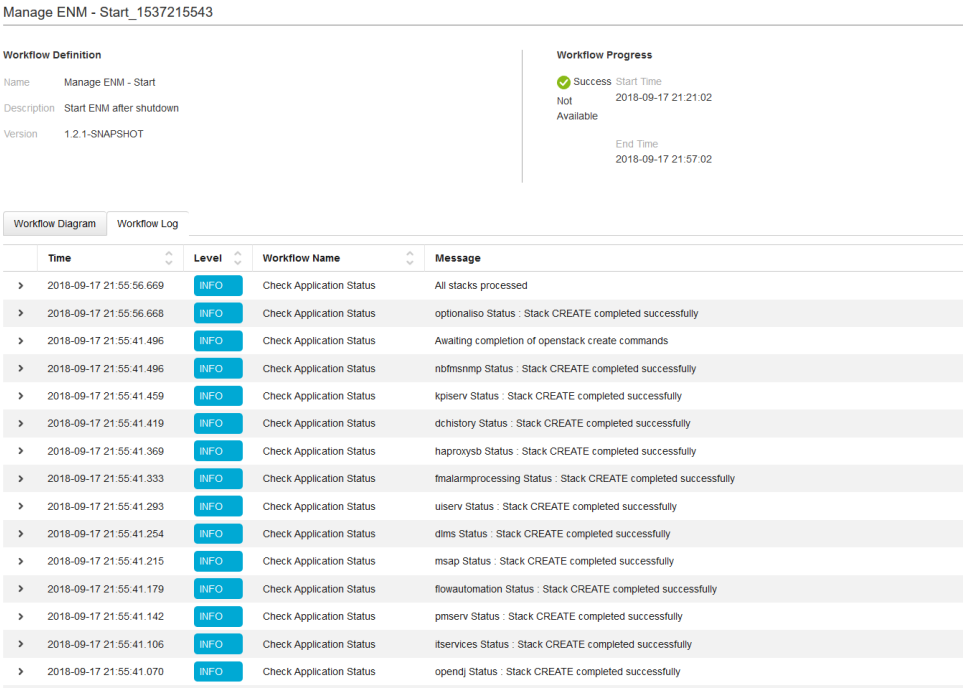

6. Gather information about the state of the system after the start.

    a. Source the OpenStack RC file from an OpenStack client machine for the ENM on Cloud deployment.

    ```
    # source <OpenStack RC file name>.rc
    ```

    b. Run the following OpenStack commands to verify the deployment state.

```
# openstack stack list
# openstack server list
# openstack volume list
```

Note:
- All stacks should be in the `CREATE_COMPLETE` state.
- All servers should be in the `ACTIVE` state.
- All volumes should be attached and in the `in-use` state.

Example

```
[root@ieatenm3314-str171 ~]# openstack stack list

+--------------------------------------+------------------------------
---+-----------------+----------------------+--------------+
| ID                                   | Stack Name                   →
   | Stack Status    | Creation Time        | Updated Time |
+--------------------------------------+------------------------------
---+-----------------+----------------------+--------------+
| fbe90028-bb7a-                       | ieatenmc7a11_wpserv_f5120752 →
   | CREATE_COMPLETE | 2018-08-23T14:31:04Z | None         |
| 4cf2-a042-64953b650536               | -651b-45ee-8a7e-82c5363824f4 →
   |                 |                      |              |
| 6f3bce85-913c-                       | ieatenmc7a11_winfiol_aba800c7 →
   | CREATE_COMPLETE | 2018-08-23T14:31:01Z | None         |
| 4e49-9067-5b6162f7a286               | -705a-408d-b403-a465ae140561 →
   |                 |                      |              |
| f355eaae-4168-4d2a-b33a-             | ieatenmc7a11_visinamingsb_ebd6 →
32 | CREATE_COMPLETE | 2018-08-23T14:30:59Z | None         |
| 56dce9c7f1d6                         | ae-31a9-4870-8c10-492eee3ac160 →
   |                 |                      |              |
 9cf6795f-3f0b-4aac-                  | ieatenmc7a11_visinamingnb_30978 →
5 | CREATE_COMPLETE | 2018-08-23T14:30:57Z | None         |
| aad8-f8d77b68038f                    | c1-8549-499c-9858-64cd74037f0e →
   |                 |                      |              |
| 94bb7eca-                            | ieatenmc7a11_said_0a6c84f7-7ef →
9  | CREATE_COMPLETE | 2018-08-23T14:30:29Z | None         |
| fb81-462c-a168-ecfb521a06d4          | -4ecf-802a-83d568946a95      →
   |                 |                      |              |
| bd88746a-d894-4f87-a76c-             | ieatenmc7a11_pmserv_a4298bfa- →
   | CREATE_COMPLETE | 2018-08-23T14:30:27Z | None         |
| 8a7941dbf8ef                         | 0d80-4fd4-97fb-df5a14315f5b  →
   |                 |                      |              |
| 1224468f-1a0b-457c-                  | ieatenmc7a11_pmrouterpolicy_b3 →
f1 | CREATE_COMPLETE | 2018-08-23T14:30:24Z | None
/.../

[root@ieatenm3314-str171 ~]# openstack server list
+--------------------------+--------------------------+--------
+--------------------------+--------------------------+
| ID                       | Name                     | Status   →
| Networks                 | Image Name               |
+--------------------------+--------------------------+--------  →
+--------------------------+--------------------------+
| 78952225-7d4c-4281-81e3-b3 | ieatenmc7a11-winfiol-0     | ACTIVE   →
| provider_network2_7A=2001: | ERICrhel6jbossimage_CXP9031 |
| b444f72172               |                          |          →
| 1b70:6207:85:0:1031:11:2c, | 560-2.53.1_CI.qcow2      |
|                          |                          |          →
| 131.160.142.174; enm_inter |                          |
|                          |                          |          →
| nal_network_Vandals_C7A11= |                          |
|                          |                          |          →
| 10.10.0.222              |                          |
| 7f4e72c9-bcd7-4239-bcd7-14 | ieatenmc7a11-winfiol-1     | ACTIVE   →
| provider_network2_7A=2001: | ERICrhel6jbossimage_CXP9031 |
| 72cf76669c               |                          |          →
| 1b70:6207:85:0:1031:11:2d, | 560-2.53.1_CI.qcow2      |
|                          |                          |          →
| 131.160.142.175; enm_inter |                          |
|                          |                          |          →
```

```
| nal_network_Vandals_C7A11= |                                |
|                            |                                |
| 10.10.0.223                |                                |
| 1a8821ff-a748-412f-        | ieatenmc7a11-wpserv-1      | ACTIVE |
| enm_internal_network_Vanda | ERICrhel6jbossimage_CXP9031 |
| 90c7-f6c882a9b627          |                                |
| ls_C7A11=10.10.0.225       | 560-2.53.1_CI.qcow2            |
| 156c041e-b160-4efc-91fd-   | ieatenmc7a11-wpserv-0      | ACTIVE |
| enm_internal_network_Vanda | ERICrhel6jbossimage_CXP9031 |
| 3b17b6e2e740               |                                |
| ls_C7A11=10.10.0.224       | 560-2.53.1_CI.qcow2            |
| 9f77908c-5b04-4b2a-        | ieatenmc7a11-visinamingsb- | ACTIVE |
| provider_network2_7A=2001: | ERICrhel6baseimage_CXP90315 |
| 96b7-7c398de752eb          | 0                              |
| 1b70:6207:85:0:1031:11:2a, | 59-2.42.1_CI.qcow2             |
|                            |                                |
| 131.160.142.172; enm_inter |                                |
|                            |                                |
| nal_network_Vandals_C7A11= |                                |
|                            |                                |
| 10.10.0.220                |                                |
/.../
[root@ieatenm3314-str171 ~]# openstack volume list
+--------------------------------------+----------------------------------
-----+--------+------+--------------------------------------+
| ID                                   | Display Name
     | Status | Size | Attached to                          |
+--------------------------------------+----------------------------------
-----+--------+------+--------------------------------------+
| 5157ca90-a974-4dd1-8220-c9aeb7152 | ieatenmc7a11-rhel7_iso_volume
-0   | in-use |    5 | Attached to ieatenmc7a11-repo-0 on  |
| 036                               |                                  |
     |        |      | /dev/vdb                             |
| 0c90c38a-2417-456a-9a0b-          | ieatenmc7a11-rhel7_updates_is
o_vo | in-use |    5 | Attached to ieatenmc7a11-repo-0 on  |
| 6820d4c72257                      | lume-0                           |
     |        |      | /dev/vdd                             |
| 6fa9d15c-b7ec-4c25-8d1b-          | ieatenmc7a11-rhel6vol-0
     | in-use |    5 | Attached to ieatenmc7a11-repo-0 on  |
| 0aac883208b1                      |                                  |
     |        |      | /dev/vdg                             |
| 43babe8a-a3d3-45c6-b20b-          | ieatenmc7a11-rhel6_updates_vo
l-0  | in-use |    5 | Attached to ieatenmc7a11-repo-0 on  |
| f4725ed3290c                      |                                  |
     |        |      | /dev/vdf
/.../
```

13 Change External NTP Server IP Addresses

This section describes a procedure to change the external NTP server IP address list for Small Integrated ENM Multi-Technology and Transport Only deployments.

Note:
- It is recommended that you change the NTP server IP address list as part of an end-to-end software upgrade. The Small Integrated ENM upgrade documents for both Transport Only and Multi-Technology deployments refer to the platform component of this procedure to apply the new NTP server IP address list to the platform at software upgrade time. The platform steps are indicated below with the prefix **Platform:**.
- If it is not possible to schedule the NTP server IP address list change for the next end-to-end software upgrade, the procedure includes a final step to run a non-software upgrade of VNF-LCM and ENM. This incurs the same downtime as a regular software upgrade.

NTP provides an essential time synchronization service. The VIO platform, VNF-LCM, and ENM components must not be without NTP service for an extended period. If possible, migrate to a new set of NTP servers that are synchronized with the existing set. This keeps the existing set of NTP servers available until the last component has switched over to the new set. If the existing NTP servers are configured with new IP addresses before starting the procedure, avoid long delays between steps to minimize the NTP service outage.

The new NTP server list is first applied to the platform components followed by VNF-LCM and ENM. All NTP servers must be synchronized and reporting the correct time.

Prerequisites

- The VNF-LCM and ENM SEDs are updated with the new NTP server IP address list and the JSON format of the updated SEDs are available. The SED setting for VNF-LCM is `<ntp_servers>`. The SED setting for ENM is `<ntp_external_servers>`.
- The new NTP server IP addresses have equivalent VLAN connectivity to the old NTP server IP addresses.
- If new NTP servers are introduced, their time is synchronized with the existing NTP servers.
- All NTP servers report the correct time.

— No VNF-LCM workflows are in progress. Check for running workflows on the VNF-LCM UI at the following URL: `http://<external_ipv4_for_services_vm>/index.html#workflows`

Steps

1. **Platform**: Change the NTP server IP address list on ESXI host(s).

    a. Log on to the vCenter client URL `https://<vcenter_ip_vio_mgt>/ui/` as `administrator@vsphere.local` with password `<vcenter_sso_password>` from the ENM SED.

    b. Select the first ESXi host in the vCenter inventory and select the **Configure** tab.

    Note: There is only one ESXi host for Transport Only deployments.

    c. Select **Time Configuration** under **System** and click the **Edit** button to open the **Edit Time Configuration** window.

    d. Check that **Use Network Time Protocol (Enable NTP client)** is selected.

    e. Check that **NTP Service Startup Policy** is set to **Start and stop with host**.

    f. Enter the new comma-separated list of NTP server IP addresses from the updated ENM SED setting `<ntp_external_servers>` for **NTP Servers**.

    g. Click **OK** to close the **Edit Time Configuration** window.

    h. On the **Configure**, tab select **Services**.

    i. Select **NTP Daemon** service and click **Restart** button.

    j. Click **OK** to confirm restart.

    k. Repeat this step for the other ESXi hosts in the cluster for Multi-Technology deployments.

2. **Platform:** Change the NTP server IP address list on vCenter Appliance (vCSA).

    a. Log on to the vCenter Server Appliance Management Interface URL `https://<vcenter_ip_vio_mgt>:5480` as `root` with password `<vcenter_sso_password>` from the ENM SED.

    b. Select **Time** in **Navigator**.

    c. Click the **Edit** button next to **Time Synchronization** to open the **Edit Time Synchronization Settings** window.

d. Check that **Mode** is set to **NTP**.

e. Enter the new comma-separated list of NTP server IP addresses from the updated ENM SED setting `<ntp_external_servers>` for **Time servers**.

f. Click **OK** to close the **Time Synchronization Settings** window.

Note: This configures the entered NTP IP list on the vCSA and restarts the NTP service

3. **Platform:** Change NTP server IP address list on VIO.

Note: It is only necessary to change the NTP server list on the VIO OMS. The OpenStack deployment nodes are already configured to sync time with the OMS.

a. Connect over SSH to the VMS on the management network IP `<vms_ip_vio_mgt>` as `root` with password `<vms_root_password>` from the ENM SED.

b. Connect over SSH to the OMS from the VMS as user `viouser` and switch to `root` user.

c. Edit the file `/etc/ntp.conf` and modify the existing NTP server IP address lines - as shown in the following example - using values from the updated ENM SED setting `<ntp_external_servers>`.

d. Save the file and restart the NTP service on the OMS as shown.

Example

```
[root@vms] # ssh viouser@oms
[viouser@oms] # sudo -i
[root@oms] # vi /etc/ntp.conf

# Modify the existing NTP server IP address lines

server <NTP server IP address 1> # First IP from the updated ENM SE →
D setting <ntp_external_servers>.
server <NTP server IP address 2> # Second IP from the updated ENM S →
ED setting <ntp_external_servers>.

# Do not change this entry
server 127.127.1.0

<save file and exit editor>

 # Restart NTP service on OMS.
[root@oms] # systemctl restart ntp

# Confirm NTP service restart.
[root@oms] # systemctl status ntp

# Confirm NTP service
[root@oms] ntpq -p
```

4. **Platform:** Change the NTP server IP address list on VMS.

a. Connect over SSH to the VMS on the management network IP `<vms_ip_vio_mgt>` as user `root` with password `<vms_root_password>` from the ENM SED.

b. Edit the file `/etc/ntp.conf` and modify the existing NTP server IP address lines - as shown in the following example - using values from the updated ENM SED setting `<ntp_external_servers>`.

c. Save the file and restart the NTP service on the VMS as shown.

Example

```
[root@vms] # vi /etc/ntp.conf

# Modify the existing NTP server IP address lines using the updated →
 ENM SED setting <ntp_external_servers>

server <NTP server IP address 1> # First IP from the updated ENM SE →
D setting <ntp_external_servers>.
server <NTP server IP address 2> # Second IP from the updated ENM S →
ED setting <ntp_external_servers>.

<save file and exit editor>

# Restart NTP service on VMS.
[root@vms] # systemctl restart ntpd

# Confirm NTP service restart.
[root@vms] # systemctl status ntpd

# Confirm NTP service
[root@vms] ntpq -p
```

5. Change the NTP server IP address list on VNF-LCM and ENM.

Note: Skip this step if the new NTP server address list is being applied as part of an end-to-end software upgrade.

— Follow document *ENM on Cloud Upgrade Instructions (2/153 72-AOM 901 151)* and perform a non-software upgrade.

Note: Ensure the software versions provided are the same as the current running ENM system..

14 Remove Orphaned VMs from vCenter Inventory

This section describes how to remove orphaned VMs from vCenter Inventory.

Prerequisites

— Healthy vSAN cluster.

— VMS VM is operational.

— OMS VM is up and running.

Steps

1. Log on to the VMS on the management network <vms_ip_vio_mgt> as user root with password <vms_root_password>.
2. Run the following command.

```
[root@vms ~]# /opt/ericsson/senm/bin/oms_bur.sh -a clean
```

Example

```
[root@vms ~]# /opt/ericsson/senm/bin/oms_bur.sh -a clean
--------------------------------------------------------------------------------
----------------
OMS Backup & Restore
--------------------------------------------------------------------------------
----------------
Enter OpenStack 'admin' user password: admin
15:52:47 INFO Cleaning out any orphaned instances
No orphaned instances found.
15:52:54 INFO Cleaning out any orphaned VMs
1 virtual machines deleted.
15:53:00 INFO Cleaning out any shadow VMs
No orphaned shadow virtual machines found.
15:53:05 INFO Log file used: /vol1/senm/log/oms_bur.sh_180605_155243.log
```

Note: If there is still an orphaned VMS VM in the vCenter UI after running the command, right-click the orphaned VM in the vCenter inventory and remove as follows: **All Virtual Infrastructure Actions** > **More Uncategorized Actions** > **Remove from Inventory**

15 VMS Artifact Cleanup Procedure

This section describes how to clean up old artifacts on VMS to free up disk space for future upgrades.

Prerequisites

ENM install or upgrade is complete.

Steps

1. Log on to VMS (using the IP address `<vms_ip_vio_mgt>`) as user `root` with password `<vms_root_password>`.
2. Run the following command to clean up the `/vol1/ENM/artifacts/` directory.

```
[root@vms ~]# ansible-playbook /opt/ericsson/senm/lib/playbooks/cleanup_ivms_artifacts.yml
```

Example

```
[root@vms ~]# ansible-playbook /opt/ericsson/senm/lib/playbooks/cleanup_ivms_artifacts.yml

PLAY [localhost] ***************************************************************
*********************************************************************************
*******************************

TASK [Remove all VMDK and QCOW2 files from '/vol1/ENM/artifacts'] **********
*********************************************************************************
*******************************
changed: [localhost] => (item=/vol1/ENM/artifacts/ERICrhel6jbossimage_CXP9031560-2.61.6.vmdk)
changed: [localhost] => (item=/vol1/ENM/artifacts/ERICrhel7baseimage_CXP9032719-1.49.1.vmdk)
changed: [localhost] => (item=/vol1/ENM/artifacts/RHEL7-Media-CXP9029081-1.0.2.vmdk)
changed: [localhost] => (item=/vol1/ENM/artifacts/ERICrhel7baseimage_CXP9032719-1.51.4.vmdk)
changed: [localhost] => (item=/vol1/ENM/artifacts/ERICrhel6baseimage_CXP9031559-2.49.4.vmdk)
changed: [localhost] => (item=/vol1/ENM/artifacts/RHEL_OS_Patch_Set_CXP9034997-2.1.2.vmdk)
changed: [localhost] => (item=/vol1/ENM/artifacts/RHEL7_OS_Patch_Set_CXP9035024-1.5.1.vmdk)
changed: [localhost] => (item=/vol1/ENM/artifacts/ERICrhelpostgresimage_CXP9032491-4.4.27.vmdk)
changed: [localhost] => (item=/vol1/ENM/artifacts/ERICrhel7baseimage_CXP9032719-1.51.4.qcow2)
changed: [localhost] => (item=/vol1/ENM/artifacts/ERICrhel6baseimage_CXP9031559-2.49.4.qcow2)
changed: [localhost] => (item=/vol1/ENM/artifacts/ERICrhelpostgresimage_CXP9032491-4.4.27.qcow2)
changed: [localhost] => (item=/vol1/ENM/artifacts/ERICrhel6jbossimage_CXP9031560-2.61.6.qcow2)
changed: [localhost] => (item=/vol1/ENM/artifacts/ERICrhel7baseimage_CXP9032719-1.49.1.qcow2)
```

```
TASK [Get all images uploaded to Glance] ********************************** →
******************************************************************************** →
*******************************
changed: [localhost]

TASK [Group images by name] *********************************************** →
******************************************************************************** →
*******************************
ok: [localhost] => (item=ERICenm_CXP9027091-1.74.74)
ok: [localhost] => (item=ERICrhel6baseimage_CXP9031559-2.49.4)
ok: [localhost] => (item=ERICrhel6jbossimage_CXP9031560-2.61.6)
ok: [localhost] => (item=ERICrhel7baseimage_CXP9032719-1.51.4)
ok: [localhost] => (item=ERICrhelpostgresimage_CXP9032491-4.4.27)
ok: [localhost] => (item=ERICrhelvnflafimage_CXP9032490-5.4.30)
ok: [localhost] => (item=RHEL6.10_Media_CXP9036772-1.0.1)
ok: [localhost] => (item=RHEL7-Media-CXP9029081-1.0.2)
ok: [localhost] => (item=RHEL7_OS_Patch_Set_CXP9035024-1.5.1)
ok: [localhost] => (item=RHEL_OS_Patch_Set_CXP9034997-2.1.2)

TASK [Get all ISO files in '/vol1/ENM/artifacts'] ************************** →
******************************************************************************** →
*******************************
ok: [localhost]

TASK [Get ISO files to delete] ******************************************** →
******************************************************************************** →
*******************************
skipping: [localhost] => (item=Media file in /vol1/ENM/artifacts: /vol1/ENM/ →
artifacts/RHEL6.10_Media_CXP9036772-1.0.1.iso)
skipping: [localhost] => (item=Media file in /vol1/ENM/artifacts: /vol1/ENM/ →
artifacts/RHEL7-Media-CXP9029081-1.0.2.iso)
skipping: [localhost] => (item=Media file in /vol1/ENM/artifacts: /vol1/ENM/ →
artifacts/RHEL7_OS_Patch_Set_CXP9035024-1.5.1.iso)
skipping: [localhost] => (item=Media file in /vol1/ENM/artifacts: /vol1/ENM/ →
artifacts/ERICenm_CXP9027091-1.74.74.iso)
ok: [localhost] => (item=Media file in /vol1/ENM/artifacts: /vol1/ENM/artifa →
cts/ERICenm_CXP9027091-1.74.23.iso)
ok: [localhost] => (item=Media file in /vol1/ENM/artifacts: /vol1/ENM/artifa →
cts/ERICenm_CXP9027091-1.69.48.iso)
ok: [localhost] => (item=Media file in /vol1/ENM/artifacts: /vol1/ENM/artifa →
cts/RHEL_OS_Patch_Set_CXP9034997-2.0.3.iso)
skipping: [localhost] => (item=Media file in /vol1/ENM/artifacts: /vol1/ENM/ →
artifacts/RHEL_OS_Patch_Set_CXP9034997-2.1.2.iso)

TASK [Remove old media files from '/vol1/ENM/artifacts'] ****************** →
******************************************************************************** →
*******************************
changed: [localhost] => (item=/vol1/ENM/artifacts/ERICenm_CXP9027091-1.74.23 →
.iso)
changed: [localhost] => (item=/vol1/ENM/artifacts/ERICenm_CXP9027091-1.69.48 →
.iso)
changed: [localhost] => (item=/vol1/ENM/artifacts/RHEL_OS_Patch_Set_CXP90349 →
97-2.0.3.iso)

PLAY RECAP **************************************************************** →
******************************************************************************** →
*******************************
localhost                  : ok=6    changed=3    unreachable=0    failed=0
```

16 Configure External Syslog Server for vSphere Log Collection

This section describes how to configure the external syslog server for vSphere log collection.

Note: The following steps are applicable to and are tested on an external syslog server that has the Red Hat Enterprise Linux 7 operating system.

A similar approach must be followed if a different operating system is installed on the external syslog server.

Prerequisites

— DVMS is installed.

— Small ENM Configuration Kit is installed.

Steps

1. SSH to DVMS as user `root` on management network `<dvms_ip_vio_mgt>` with `<dvms_root_password>`.
2. Transfer the appropriate rsyslog configuration template on to the external syslog server.
   — If the deployment is Small Integrated ENM Transport Only, transfer `vio_syslog_single_host.conf_template /etc/rsyslog.d/vio_syslog_single_host.conf` to the external syslog server.

   ```
   # scp /opt/ericsson/senm/templates/vio_syslog_single_host.conf_template →
    <external_syslog_ip_vio>:/etc/rsyslog.d/vio_syslog_single_host.conf
   ```

   — If the deployment is Small Integrated ENM Multi-Technology, transfer `vio_syslog_multiple_hosts.conf_template /etc/rsyslog.d/vio_syslog_multiple_hosts.conf` to the external syslog server.

   ```
   # scp /opt/ericsson/senm/templates/vio_syslog_multiple_hosts.conf_templ →
   ate <external_syslog_ip_vio>:/etc/rsyslog.d/vio_syslog_multiple_hosts.conf
   ```

   Note: Where `<external_syslog_ip_vio>` is the IP address of the external syslog server.
3. Transfer the log rotation configuration template `/etc/logrotate.d/external_syslog_logrotation` to the external syslog server.

```
# scp /opt/ericsson/senm/templates/external_syslog_logrotation_template <ext
ernal_syslog_ip_vio>:/etc/logrotate.d/external_syslog_logrotation
```

Note: Where `<external_syslog_ip_vio>` is the IP address of the external syslog server.

4. Connect over SSH on to the external syslog server with `<external_syslog_ip_vio>`.
5. Create the log directory on to the external syslog server (this log directory will be used to collect the vSphere logs).
6. Change the content type of the log directory.

```
# chcon -R -t var_log_t <log_directory>
```

Note: Where `<log_directory>` is an absolute path of the directory created to store the vSphere logs.

7. Update the rsyslog configuration file for ESXi host information and the log directory to collect the vSphere logs.
   - If the deployment is Small Integrated ENM Transport Only, update the value of the following placeholders in the file `/etc/rsyslog.d/vio_syslog_single_host.conf`

Placeholder	Description
`<esxi_host1_ip_vio_mgt>`	IP address on the VIO Management Network for ESXi host 1. It is a SED parameter.
`<esxi_host1_vio_mgt_hostname>`	Hostname on the VIO Management Network for ESXi host 1. It is a SED parameter.
`<log_directory>`	It is an absolute path of the directory created to store the vSphere logs.

   - If the deployment is Small Integrated ENM Multi-Technology, update the value of the following placeholders in the file`/etc/rsyslog.d/vio_syslog_multiple_hosts.conf`

Placeholder	Description
`<esxi_host1_ip_vio_mgt>`	IP address on the VIO Management Network for ESXi host 1. It is a SED parameter.
`<esxi_host1_vio_mgt_hostname>`	Hostname on the VIO Management Network for ESXi host 1. It is a SED parameter.

Placeholder	Description
<esxi_host2_ip_vio_mgt>	IP address on the VIO Management Network for ESXi host 2. It is a SED parameter.
<esxi_host2_vio_mgt_hostname>	Hostname on the VIO Management Network for ESXi host 2. It is a SED parameter.
<esxi_host3_ip_vio_mgt>	IP address on the VIO Management Network for ESXi host 3. It is a SED parameter.
<esxi_host3_vio_mgt_hostname>	Hostname on the VIO Management Network for ESXi host 3. It is a SED parameter.
<log_directory>	It is an absolute path of the directory created to store the vSphere logs.

8. Restart the rsyslog service.

```
# systemctl restart rsyslog
```

9. Update the <log_directory> information in the file /etc/logrotate.d/external_syslog_logrotation.

Note:
- Where <log_directory> is the absolute path of the directory created to store the vSphere logs.
- The collection of vSphere logs begins once the **Initial Install** or the **Upgrade** workflow for a Small Integrated ENM deployment is completed. Regular disk usage is performed on the <log_directory> of the external syslog server. Where <log_directory> is an absolute path of the directory created to store the vSphere logs.

17 Update vCenter Passwords

> The following sections detail how to update the vCenter administrator or user passwords.

17.1 Update the vCenter Administrator Password

> The following section details how to update the vCenter administrator password.

Note: The vCenter Single Sign-On administrator password never expires. It is not necessary to change the administrator password if the vSphere banner message `Your password will expire in <nn> days` is seen. This message can be ignored and once the false expiration date passes, the banner will not appear any more.

Steps

1. Access the vCenter client at the following address.

   ```
   https://<vcenter_ip_vio_mgt>/ui/
   ```

2. Log in with administrator privileges as user `administrator@vsphere.local` with your current password.
3. Navigate to **Menu > Administration > Single Sign-On > Users and Groups**.
4. Select **vsphere.local** from the drop-down in **Domain**.
5. Select user appropriate vCenter user in the right side and click on three vertical dots and select **Edit**.
6. Enter **Password** and confirm **Confirm Password,** and then click **SAVE**.

   Note: Refer to the SED for password requirements.

7. Click **OK**.
8. Navigate to **Menu > VMware Integrated OpenStack**.
9. Select **OpenStack Deployments**.
10. Click the **Manage** option.
11. Click the **Settings** option.

12. Click **Change Password**.
13. Select and complete the following in the hostname from the drop-down menu.
    - **vCenter IP address:** <vcenter_ip_vio_mgt>
    - **Username:** administrator@vsphere.local
    - **New Password:** must be the same as the password entered at step 6.
    - **Confirm Password:** must be the same the password entered at step 6.
14. Click **Submit** and **OK** when the warning about synchronizing the vCenter username and password is displayed.
15. Wait until the OpenStack configuration is complete and it is in a running state again.
16. Log on to the VMS on the management network <vms_ip_vio_mgt> as user root with password <vms_root_password>.
17. Update the new vCenter administrator password <vcenter_sso_password> in the *Site Engineering Data for ENM on Cloud (2/1057-AOM 901 151)* and update <vcenter_sso_password> in /vol1/senm/etc/sed.json.
18. Update the govc.rc file.

```
[root@vms ~]# mv /vol1/senm/etc/govc.rc /vol1/senm/etc/govc.rc.backup
[root@vms ~]# /opt/ericsson/edpcore/bin/edp_autodeploy.sh -p sienm_post_inst →
all_i -l
[root@vms ~]# source /vol1/senm/etc/govc.rc
```

19. Log on to OMS as viouser and switch to root.

```
[root@vms ~]# ssh viouser@oms
viouser@ieatvio028-vio-mgt-4:~$ sudo -i
```

20. Check VIO status, all services must be in a running state.

    Note: It can take up to 10 minutes for all services to be in a running state.

```
root@ieatvio028-vio-mgt-4:~# viocli deployment status --period 60
Collector Name                      Overall Status
----------------------------------  ----------------
VerifyRabbitmqMemory                SUCCESS
VerifyTimeSynchronization           SUCCESS
VerifyDatabaseConnectionPerProcess  SUCCESS
VerifyRunningProcess                SUCCESS
VerifyConnection                    SUCCESS
VerifyRabbitmqDescriptor            SUCCESS
VerifyMariaDatabaseClusterSize      SUCCESS
```

17.2 Update Other vCenter User Passwords

> The following section details how to update the password associated with the vCenter users `vcenter_backup_username`, `vcenter_monitor_username`, and `vcenter_drs_username`.

Note: `vcenter_drs_username` is valid for Multi-Technology deployments only.

Prerequisites

ENM on Cloud SED is available at `/vol1/senm/etc/sed.json` on VMS.

Steps

1. Access the vCenter client at the following address.

   ```
   https://<vcenter_ip_vio_mgt>/ui/
   ```

2. Log in with administrator privileges as user `administrator@vsphere.local` and password `<vcenter_sso_password>`.
3. Navigate to **Menu > Administration > Single Sign-On > Users and Groups**.
4. Select **vSphere.local** from the **Domain** drop-down.
5. Select an appropriate vCenter user in the right-side and select **Edit**.
6. Enter **Password**, **Confirm Password** and click **SAVE**.

   Note: Refer to the SED for password requirements.

7. Log on to the VMS on the management network `<vms_ip_vio_mgt>` as user `root` with password `<vms_root_password>`.
8. Update the password values (`vcenter_backup_username`, `vcenter_monitor_username`, `vcenter_drs_username`) as needed in the `/vol1/senm/etc/sed.json` file.
9. Update the consul with the new passwords.

   ```
   [root@vms ~]# /opt/ericsson/senm/bin/configure_vms.sh -a addcreds -e /vol1/s →
   enm/etc/sed.json
   ```

18 Configure Licenses

This section describes how to apply or update VMware licenses.

Stop!

This step is only applicable if VIO is upgraded to a major VIO version during completion of the section *Post-Upgrade Configuration and Cleanup* in ENM on Cloud Upgrade Instructions.

Prerequisites

- VMware Integrated OpenStack (VIO) platform is installed and is in a healthy state.
- Licenses are available.
- ENM SED is available on VMS VM.
- VNF-LCM SED is available on VMS VM.

Steps

1. Remove the old VIO solution from vCenter.

   Note: This step is only applicable if VIO is upgraded to a major VIO version during completion of the section *Post-Upgrade Configuration and Cleanup* in ENM on Cloud Upgrade Instructions .

   a. Log on to the vCenter Client URL `https://<vcenter_ip_vio_mgt>/ui/` as administrator `administrator@vsphere.local` with password `<vcenter_sso_password>`.

   b. Click on the **Menu** > **Administration**.

   c. In the left panel, select **Licenses** menu icon at the top of the under **Licensing**.

   d. Select **Assets** and then click on the **Solutions** tab.

   e. Select the check box beside the older version of the VIO solution and select **X Remove Solution**.

   f. Select **Yes** in the `Remove Solution` warning popup that appears.

2. Log on to the VMS on the management network <vms_ip_vio_mgt> as user root with password <vms_root_password>.

3. Create the licenses.yml file by copying the license file template from its repository to the /vol1/senm/etc/ directory.

```
[root@vms ~]# cp /opt/ericsson/senm/templates/licenses.yml_template /vol1/se
nm/etc/licenses.yml
```

4. Open the licenses.yml file for editing and enter the license key for each listed product.

```
[root@vms ~]# vim /vol1/senm/etc/licenses.yml
```

5. Run the edp_autodeploy.sh script with the sienm_licenses_i profile to add the license keys to vCenter and assign them to their respective products.

```
[root@vms ~]# /opt/ericsson/edpcore/bin/edp_autodeploy.sh -e /vol1/senm/etc/
sed.json -m /vol1/senm/etc/lcm_sed.json -p sienm_licenses_i
```

Example

```
[root@vms ~]# /opt/ericsson/edpcore/bin/edp_autodeploy.sh -e /vol1/senm/etc/
sed.json -m /vol1/senm/etc/lcm_sed.json -p sienm_licenses_i
--------------------------------------------------------------------------
-----------------
EDP Autodeploy vms
--------------------------------------------------------------------------
-----------------
09:11:24 INFO Blank Ansible answer file Created [/vol1/senm/etc/answerfile.y
ml]
2020-02-18 09:11:24,878 main INFO Running setup and configuration for select
ed profiles.
2020-02-18 09:11:24,881 get_product_abbreviations INFO All profile abbreviat
ions being used: ['sienm']
09:11:31 INFO Ansible deploy SED Generated [/vol1/senm/etc/deploy_sed.yml]
09:11:32 INFO Ansible answerfile Updated [/vol1/senm/etc/answerfile.yml]
2020-02-18 09:11:32,058 execute_scripts_in_directory INFO Executed /opt/eric
sson/edp/automation/sienm/config/setup_and_read_sed.sh successfully.
2020-02-18 09:11:32,059 main INFO Setup and configuration completed successf
ully...
2020-02-18 09:11:32,168 main INFO Command line arguments validation complete
d successfully...
OK to start at stage sienm_apply_licenses.yml? [YeS|n] (n): YeS
09:11:56 INFO 'YeS' selected
09:11:56 INFO Run stage: ----> sienm_apply_licenses.yml
09:12:07 INFO Log file used: vms:/vol1/senm/log/edp_autodeploy.sh_200218_091
124.log
```

6. Back up the licenses.yml file to a secure location once the licenses are successfully applied and delete it from VMS for security reasons.

```
[root@vms ~]# rm -f /vol1/senm/etc/licenses.yml
```

19 vCenter Alarm Whitelist

This section describes a list of alarms in vCenter that can be ignored during operating.

Table 5 vCenter Alarm Whitelist

Alarm Name	Description
vSAN health alarm `vSAN disk balance`	This means that the disk load variance between some disks exceeded the threshold at some point in time. vSAN will re-actively load balance the cluster. This alarm can be safely ignored.
`Insufficient configured resources to satisfy the desired vSphere HA failover level on the cluster <vcenter_object_prefix>_CLUS in <vcenter_object_prefix>_DC`	VMware GSS confirmed this is a bug that will be addressed in the next patch release. This alarm is on when enable override in Admission Control. This alarm can be safely ignored.
The device state displays one of the following warnings: `[Device] I/O Module <number> ALOM_Link_P<number>` or `[Device] I/O module <number> NIC_Link_P<number>`	This issue is reported with HPE ProLiant servers running VMware ESXi 6.0, VMware ESXi 6.5, or VMware ESXi 6.7 with HPE Integrated Lights-Out 5 (iLO 5) firmware version 1.30.
vSAN health alarm `Disk format version`	This means that the disk format version of one or more vSAN disks are compatible, but out of date. This is observed after updating to vSphere 6.7. This alarm can be safely ignored and alarm can be **Reset to Green**.
Registration/unregistration of third-party IO filter storage providers fails on a host	This issue is caused by a transient connectivity issue between the ESXi hosts and the vCenter Server while the host is being provisioned in the vSphere cluster. This alarm can be safely ignored and alarm can be **Reset to Green**.

20 Troubleshooting

> This document describes the troubleshooting steps and known issues for a Small Integrated ENM deployment.

20.1 Ansible Troubleshooting

> This section outlines general steps to troubleshoot a failed Ansible playbook.

This section is intended as a general guideline. No specific fault is identified in this section.

Note: When `ansible-playbook` fails, it can be rerun with additional flags.

Ansible Logs

Ansible playbook executions are logged to `/vol1/senm/log/ansible.log`.

Syntax Check

After any updates of the ENV vars file or playbook, run first with the `--syntax-check` flag.

Note: This will not execute any commands, only verify the input files.

```
# ansible-playbook -e "@<env>" <playbook> --syntax-check
```

Verbosity

After a failed execution, the verbosity can be increased for a second run.

Note: The verbosity level can be selected in four levels.

```
# ansible-playbook -e "@<env>" <playbook> -v[v[v[v]]]
```

Retry

To only rerun the failed task, a separate playbook has been prepared by Ansible and can be launched as follows.

```
# ansible-playbook -e "@<env>" <playbook>.retry
```

Note: The playbook should be idempotent, so there is actually no problem rerunning all tasks from the beginning, but you save time and output, if you have on verbose flag.

Start at Specific Task

First list the names of all tasks in the playbook, and then specify the desired task name where to start execution.

```
# ansible-playbook -e "@<env>" <playbook> --list_tasks
# ansible-playbook -e "@<env>" <playbook> --start-at-task="<task name>"
```

Steps

To run each task step-by-step, with a prompt between each task, use the `--step` flag.

```
# ansible-playbook -e "@<env>" <playbook> --step
```

Note: This flag can be combined with the `--start-at-task` and `-vvvv` flags.

Tune Variables

Variables can be set on the command line. You can either add a new variable or override an existing variable in the ENV file.

```
# ansible-playbook -e "@<env>" <playbook> --extra-vars <my-var>="<new value>"
```

Variables read from the environment by the playbook, can be set in the shell like this.

```
# export my-var="MyValue"
# ansible-playbook -e "@<env>" <playbook>
In playbook:
  var_x: "{{ lookup('env', 'my-var') }}"
```

Debug Variable

Variables can be examined during execution. To print a variables value, use the Ansible module `debug`. The playbook has to be edited.

```
# vi <playbook>
```

Add this section at the desired position in the playbook, for example between to existing tasks. Replace *<my-var>* with the desired variable name.

```
  - debug: msg="{{ <my-var> }}"
```

Debug Python Code

Most modules are written in Python and when they fail it is difficult to debug and trace.

The code is run on a remote host (or local), from a temporary directory, which is removed immediately after failure, before exit.

Set an environment variable in the shell, and rerun the playbook. Either add `-v` to the bash script or `-vvvv` to the ansible-playbook command.

```
# export ANSIBLE_KEEP_REMOTE_FILES=1
# configure_vms.sh -a <action> -e <sed> -v
OR:
# ansible-playbook -e "@<env>" <playbook> -vvvv
```

From the output, copy-paste the line running the python script, and run it with some additional parameters. First `explode` then `execute`:

Example 1

```
. . . in the output from above:
EXEC /bin/sh -c '/usr/bin/python2 /root/.ansible/tmp/ansible-tmp-1522179483.74-2 →
50289391849642/firewalld.py && sleep 0'
. . . now do:

# /root/.ansible/tmp/ansible-tmp-1522179483.74-250289391849642/firewalld.py expl →
ode
# cd /root/.ansible/tmp/ansible-tmp-1522179483.74-250289391849642/debug_dir
# ls -l
# vim ansible_module_firewalld.py
# /root/.ansible/tmp/ansible-tmp-1522179483.74-250289391849642/firewalld.py exec →
ute
```

20.2 Resolve Certificate Error

> This section describes the steps to resolve a certificate error that may be encountered when uploading files to a datastore on VCSA.

Steps

Note: If you encounter the error shown in the example image, follow these steps.

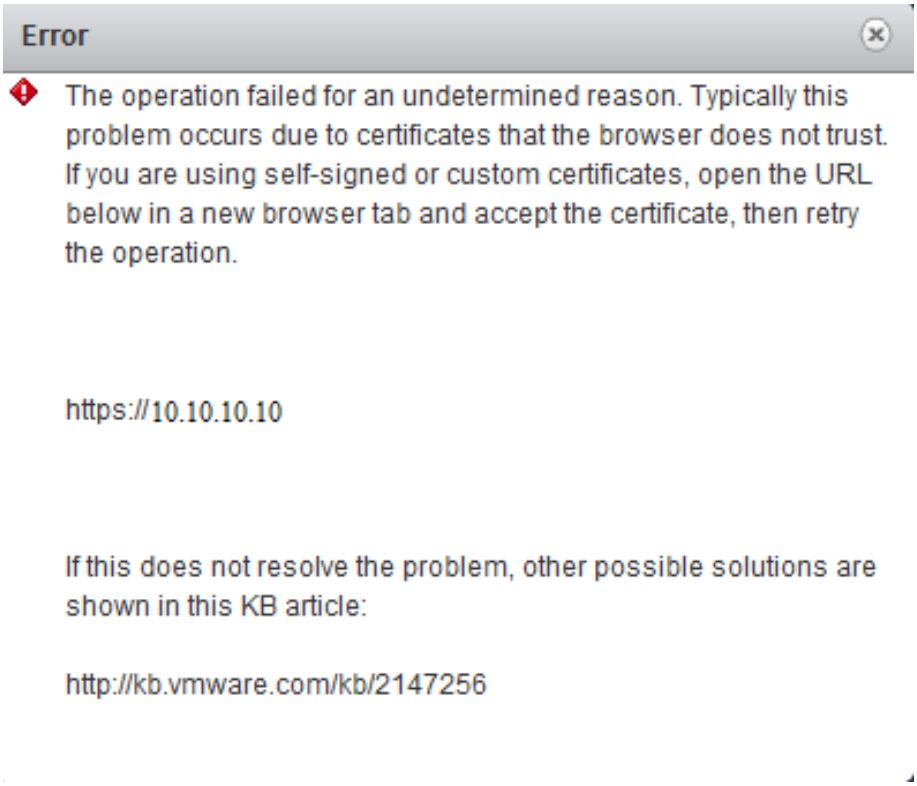

1. Open the first URL shown in the error message, typically the IP address for an ESXi host.
2. Resolve the certificate error by the applicable method for your browser.
   - For Google Chrome, click **Show Advanced**, then **Proceed to <x>** where <x> is the URL shown in the error message.
   - For Mozilla Firefox, click **I Understand the Risks, Add Exception**, then Confirm Security Exception.
   - For Internet Explorer, click **Continue to this website (not recommended)**.

   Note: The error is now resolved and you may proceed with the upload.

Results

The certificate error has been resolved.

20.3 Recover VM in PXE Boot Mode

> This section provides steps to recover an instance VM that goes to PXE boot instead of booting from the internal boot disk.

If an instance VM reboot results in a PXE boot with error `Operating System not found` instead of booting from the internal boot disk, run the following steps to recover the VM boot process.

Prerequisites

The VM tries to PXE boot and fails with console error `Operating System not found`.

Steps

1. Log on to vCenter, right-click the VM and select **Power** > **Power Off**.
2. Right-click the powered-off VM and select **Edit Settings**.
3. Select the **VM Options** tab and select **Boot Options > Force BIOS setup** and click **OK** to save the VM settings.
4. Power on the VM and select the **Open Console** option to display the BIOS settings menu.
5. Select the **Boot** tab on the BIOS menu,and use the **+** and **-** keys to move the boot disk, normally **Hard disk 1**, to the top of the boot order list.
6. Press **F10** to save the settings.
7. Boot the instance with the following command on VMS.

```
root@vms # openstack server start <server_name>
```

Note: Where `<server_name>` is the OpenStack server name.

20.4 Known Issues and Troubleshooting a VIO Install

> The following section can be used to troubleshoot a VMware Integrated OpenStack(VIO) installation.

20.4.1 VMware Integrated OpenStack Installation Fails

> Troubleshooting an installation failure of VMware Integrated OpenStack(VIO).

Cause: Installation of VIO fails because of time out.

The issue occurs when deploying OpenStack from vCenter through the VIO plug-in.

The progress bar gets to 75% and some time later a time-out error displays on the GUI.

The following logs were gathered from the loadbalancer around the time of the failure.

```
Nov 29 15:46:45 loadbalancer01 ansible-sysctl: Invoked with name=net.ipv4.ip_loc →
al_port_range ignoreerrors=False value=10000 65535 reload=True state=present sys →
ctl_set=True sysctl_file=/etc/sysctl.conf
Nov 29 15:47:00 loadbalancer01 rsyslogd-2207: error during parsing file /etc/rsy →
slog.d/49-haproxy.conf, on or before line 2: warnings occured in file '/etc/rsys →
log.d/49-haproxy.conf' around line 2 [v8.16.0 tryhttp://www.rsyslog.com/e/2207 ]
Nov 29 15:47:38 loadbalancer01 Keepalived_vrrp[947]: Netlink: filter function er →
ror
Nov 29 15:47:38 loadbalancer01 Keepalived_vrrp[947]: Netlink: filter function er →
rorAfter a number (2 - 3) retries, installation will succeed.
```

Solution

1. After a number of retries (2 to 3), installation will succeed.

20.4.2 Failed to SSH to iLO during Installation of Small Integrated ENM Transport Only

> Installation of VIO fails because of iLO SSH connection failure.

Diagnostics

The issue can occur when deploying VIO at stage `prepare_kickstart`.

The following error message was gathered from log file `/vol1/senm/log/ilo_mount.log`

```
TASK [Mount new ISO as virtual CD] ******************************************** →
********************************************************************************  →
fatal: [localhost]: FAILED! => {"changed": true, "cmd": "/opt/ericsson/senm/lib/  →
ilo_mount_cd.exp 10.151.11.7 root shroot12 http://159.107.161.166/custom_ESXi.is →
o >/vol1/senm/log/ilo_mount.log", "delta": "0:00:02.079301", "end": "2018-08-28  →
12:02:47.528484", "msg": "non-zero return code", "rc": 255, "start": "2018-08-28 →
 12:02:45.449183", "stderr": "", "stderr_lines": [], "stdout": "", "stdout_lines →
": []}
```

1. Log on to <esxi_host1_ip_ilo> as user <esxi_host1_ilo_user> with password <esxi_host1_ilo_password>.
2. Reset the iLO using following command:

   ```
   </>hpiLO-> reset /map1
   ```

   Example

   ```
   </>hpiLO-> reset /map1

   status=0
   status_tag=COMMAND COMPLETED
   Fri Sep  7 09:30:46 2018

   Resetting iLO.

   CLI session stopped
   Received disconnect from 10.151.41.226 port 22:11:  Client Disconnect
   Disconnected from 10.151.41.226 port 22
   ```

20.5 Known Issues and Troubleshooting for ENM Upgrade on VIO

> Troubleshooting an ENM upgrade on VMware Integrated OpenStack(VIO).

Refer to the topics contained in the sections *Troubleshooting the Deployment* and *Rollback ENM on Cloud* in *ENM on Cloud Upgrade Instructions (2/153 72-AOM 901 151)*.

Note: The following caveats apply to Small Integrated ENM deployments.

Caveats

— The section *Prevent MTU Size Defaulting on VNF-LCM VMs* is not applicable to Small Integrated ENM deployments and must not be run.

— The section *Workaround to Prevent Broken Pipe Issue* in ENM ISO prior to 18.02 is not applicable to Small Integrated ENM deployments and must not be run.

— The section *Snapshot Workflow Failure - Delete Cloned Volumes* has additional steps which need to be run.

Note: Refer to Delete Orphaned VM Snapshots on VIO (Method 1) on page 120.

— All references to the `keystonerc` file refer to the `project.rc` file which can be found at `/vol1/senm/etc/<vio_os_project_name>_project.rc`

— The following table should be used to determine which properties from the SED should be used in place of any reference to `<CLOUD MANAGER IP>`, `<CLOUD MANAGER HOSTNAME>`, or `<PORT NUMBER OF KEYSTONE SERVICE>`.

Parameter	SED Parameter	Comment
`<CLOUD MANAGER IP>`	`<vio_deploy_ip_api>`	
`<CLOUD MANAGER HOSTNAME>`	`<vio_deploy_api_hostname>`	
`<PORT NUMBER OF KEYSTONE SERVICE>`	`cloudManagerRestInterfaceBaseURL`	Use the port number from this URL. Normally `5000`.

20.5.1 Delete Orphaned VM Snapshots on VIO (Method 1)

> This section describes the steps required to delete orphaned VMs immediately after a **Snapshot deployment** workflow has been canceled.

Steps

1. Check for flag `ddb.deletable = "false"`.
    a. Log on to the VMS as user **`root`**.
    b. From the VMS, SSH to the ESXi host with IP **`<esxi_host1_ip_vio_mgt>`** as user root with password **`<esxi_host1_mgt_password>`**.
    c. Change directory to the volume datastore.
        — For a SIENM Multi-Technology deployment:

        ```
        [root@esxi_1]# cd /vmfs/volumes/vsanDatastore
        ```

        — For a SIENM Transport Only deployment:

        ```
        [root@esxi_1]# cd /vmfs/volumes/datastore1
        ```

    d. Search for flag `ddb.deletable = "false"` in volume VMDK files as described below:

```
[root@esxi_1]# grep -ir "ddb.deletable.*false" */*[\)].vmdk 2>/dev/
null | grep -iv "image" | grep -iv ERICvms
```

Contact Ericsson support for the issue recovery if search results are found.

Note: Do not continue this procedure until after support have handled the flag.

2. Log on to VNF-LCM Services as `cloud-user`.

```
# ssh -i /vol1/senm/etc/key_pair_<deployment_id>.pem cloud-user@<VNFLCM-services
 IP>
```

3. Switch to `root` user.

```
# sudo -i
```

4. Run the following commands to delete VM snapshots.

```
# cd /opt/ericsson/ERICenmdeploymentworkflowsworkflows/scripts/enmdeployment
workflows/<latest workflow version>/vio/
# python clone_snap_delete.py -s <vcenter_ip_vio_mgt> -u Administrator@vsphere
.local -S -n <snapshot tag> -d all
```

Note: The `<snapshot tag>` is the tag used when starting the snapshot workflow.

5. Confirm cleanup after a canceled backup.

    a. Log on to the VMS as user `root`.

    b. From the VMS, SSH to the ESXi host with IP `<esxi_host1_ip_vio_mgt>` as user `root` with password `<esxi_host1_mgt_password>`.

    c. Change directory to the volume `datastore`.

    For a SIENM Multi-Technology deployment.

    ```
    [root@esxi_1]# cd /vmfs/volumes/vsanDatastore
    ```

    For a SIENM Transport Only deployment.

    ```
    [root@esxi_1]# cd /vmfs/volumes/datastore1
    ```

    d. Search for flag `ddb.deletable = "false"` in volume VMDK files as described below:

    ```
    [root@esxi_1]# grep -ir "ddb.deletable.*false" */*[\)].vmdk 2>/dev/
    null | grep -iv "image" | grep -iv ERICvms
    ```

Contact Ericsson support for the issue recovery if search results are found.

e. Search for snapshot delta disks.

```
[root@esxi_1]# find . -name "*-0000*.vmdk" | grep -iv "image" | gre →
p -iv "template" | grep -iv ERICvms
[root@esxi_1]#
```

Contact Ericsson support for the issue recovery if search results are found.

6. Log on to the VMS on the management network <vms_ip_vio_mgt> as root user with password <vms_root_password>.

7. Check that there are no volumes with the <snapshot tag>.

```
[root@vms log]# openstack volume list |grep <snapshot tag>
```

Note: Where <snapshot tag> is the tag used when starting the snapshot workflow.

8. Delete any snapshot volume that exists.

```
[root@vms log]# openstack volume delete <Volume ID>
```

9. Check that there are no orphaned volumes left in the data store.

```
[root@vms log]# govc datastore.ls -l ./<snapshot tag>*
```

10. Delete any orphaned volumes left in the data store.

```
[root@vms log]# govc datastore.rm  ./<snapshot volume>
```

Note: The orphaned volumes must be deleted individually, wildcards are not supported.

Results

Deployment is ready to rerun the snapshot workflow.

20.5.2 Delete Orphaned VM Snapshots on VIO (Method 2)

> This section provides the steps required to delete orphaned VM snapshots that were not cleaned up immediately after a canceled **Snapshot deployment** workflow.

Stop!

Only run the following procedure if there are no instances of **Backup Deployment** or **Snapshot Deployment** running on VNF-LCM.

Prerequisites

ENM has been fully installed.

Steps

1. Check for flag `ddb.deletable = "false"`.

    a. Log on to the VMS as user **`root`**.

    b. From the VMS, SSH to the ESXi host with IP **`<esxi_host1_ip_vio_mgt>`** as user root with password **`<esxi_host1_mgt_password>`**.

    c. Change directory to the volume datastore.

    — For a SIENM Multi-Technology deployment:

    ```
    [root@esxi_1]# cd /vmfs/volumes/vsanDatastore
    ```

    — For a SIENM Transport Only deployment:

    ```
    [root@esxi_1]# cd /vmfs/volumes/datastore1
    ```

    d. Search for flag `ddb.deletable = "false"` in volume VMDK files as described below:

    ```
    [root@esxi_1]# grep -ir "ddb.deletable.*false" */*[\)].vmdk 2>/dev/ →
    null | grep -iv "image" | grep -iv ERICvms
    ```

    Contact Ericsson support for the issue recovery if search results are found.

    Note: Do not continue this procedure until after support have handled the flag.

2. Do the following steps:

a. Log on to VMS.

b. Run a platform health check.

```
[root@vms ~]# /opt/ericsson/senm/bin/sienm_hc.sh
```

c. View the health check report.

```
[root@vms ~]# elinks /vol1/senm/log/html/HC_Report_<timestamp>.html
```

d. Check if any VMs have snapshots under **Orphaned Resources**.

3. Log on to the vCenter client at `https://<vcenter_ip_vio_mgt>/ui/` as user `<vcenter_monitor_user>` with password `<vcenter_monitor_password>` from the SED.

4. Expand the datacenter **<vcenter_object_prefix>_DC** under **Navigator** and select the cluster inventory object **<vcenter_object_prefix>_CLUS** in the left pane.

5. Select the **VMs** tab in the right pane (all VMs in the deployment are listed).

6. Select each of the VMs identified in the platform health check.

   Note: Multiple VMs can be selected by selecting the first VM, holding down SHIFT, and selecting the last VM.

7. Right-click on the selected VMs and select **Snapshots > Delete All Snapshots** from the context menu.

8. Verify that all VM snapshots have been deleted.

   Note: You may need to refresh the view. It may also be necessary in rare cases to repeat the delete snapshots operation.

9. Confirm cleanup after a canceled backup.

   a. Log on to the VMS as user `root`.

   b. From the VMS, SSH to the ESXi host with IP `<esxi_host1_ip_vio_mgt>` as user `root` with password `<esxi_host1_mgt_password>`.

   c. Change directory to the volume `datastore`.

   For a SIENM Multi-Technology deployment.

```
[root@esxi_1]# cd /vmfs/volumes/vsanDatastore
```

   For a SIENM Transport Only deployment.

```
[root@esxi_1]# cd /vmfs/volumes/datastore1
```

d. Search for flag `ddb.deletable = "false"` in volume VMDK files as described below:

```
[root@esxi_1]# grep -ir "ddb.deletable.*false" */*[\)].vmdk 2>/dev/ →
null | grep -iv "image" | grep -iv ERICvms
```

Contact Ericsson support for the issue recovery if search results are found.

e. Search for snapshot delta disks.

```
[root@esxi_1]# find . -name "*-0000*.vmdk" | grep -iv "image" | gre →
p -iv "template" | grep -iv ERICvms
[root@esxi_1]#
```

Contact Ericsson support for the issue recovery if search results are found.

10. Log on to the VMS on the management network `<vms_ip_vio_mgt>` as `root` user with password `<vms_root_password>`.

11. Check that there are no volumes with the `<snapshot tag>`.

```
[root@vms log]# openstack volume list |grep <snapshot tag>
```

Note: Where `<snapshot tag>` is the tag used when starting the snapshot workflow.

12. Delete any snapshot volume that exists.

```
[root@vms log]# openstack volume delete <Volume ID>
```

13. Check that there are no orphaned volumes left in the data store.

```
[root@vms log]# govc datastore.ls -l ./<snapshot tag>*
```

14. Delete any orphaned volumes left in the data store.

```
[root@vms log]# govc datastore.rm  ./<snapshot volume>
```

Note: The orphaned volumes must be deleted individually, wildcards are not supported.

Results

All the remaining orphaned VM snapshots have been deleted successfully.

20.5.3 ENM Upgrade Workflow Error

> This section describes steps to recover from known issues with the ENM upgrade workflow for Small Integrated ENM.

— The graphical workflow on the VNF-LCM web GUI indicates that the workflow is hung or failed.

— Check the following log file on VNF-LCM Services for error logs relating to the workflow: `/ericsson/3pp/jboss/standalone/server.log`.

Cause: Diagnostics

The ENM upgrade workflow hangs. An error similar to the following appears in the upgrade workflow log accessible over the VNF-LCM Services GUI:

```
Stack update failed for :vio-5587-elasticsearch-0 reason:
ConnectFailure: resources.elasticsearch_definition_volume_attach:
Unable to establish connection to...('Connection aborted.',
BadStatusLine("".))
```

Solution

1. Log on to the VMS as `root` user.
2. Source the project RC file.

```
[root@vms ~]# source /vol1/senm/etc/vio_<deployment name>.rc
```

3. List the server in the upgrade workflow error, in this example `elasticsearch-0` (there may be duplicate servers returned.)

```
[root@vms ~]# openstack server list | grep -i elasticsearch-0
[root@vms ~]# openstack server list | grep elastic
| ac5c1875-bfde-4c4a-afd5-cdb7659e9ed5 | vio-5578-elasticsearch-0        | A →
CTIVE | vio_internal_5578=10.10.2.19                                    →
                                                | ERICrhel6baseimage_CXP903155 →
9-2.40.3    |
| 22531276-604f-49b2-b0aa-e005e83c1f70 | vio-5578-elasticsearch-0        | A →
CTIVE |                                                                 →
                                                | ERICrhel6baseimage_CXP903155 →
9-2.40.3    |
[root@vms ~]#
```

4. Delete both servers from OpenStack using their server ID (the upgrade workflow continues and recreates or upgrades the server on the next attempt).

```
[root@vms ~]# openstack server delete ac5c1875-bfde-4c4a-afd5-cdb7659e9ed5
[root@vms ~]# openstack server delete 22531276-604f-49b2-b0aa-e005e83c1f70
```

20.5.4 Failed to Attach Volume to Server after Failed Rollback

> This section shows how to recover an ENM system if a volume has failed to attach to a server.

Prerequisites

— A rollback has failed because of a volume failing to attach to a server.

— ENM has been installed.

Steps

1. Cancel any existing HA workflows on the VNF-LCM UI.

    a. Open the VNF-LCM UI using the below URL in the browser.

    ```
    [root@vnflaf-services ~http://<external_ipv4_for_services_vm>/index.htm →
    l#workflows
    ```

    Note: Replace the value for `<external_ipv4_for_services_vm>` in the URL with the IP address from the VNF-LCM SED. If authentication is requested, log in using the username `vnfuser` and the password that was defined during the install.

    b. Look for HA workflows that are ongoing under **Instance Activity** in the right side of the VNF-LCM UI.

    c. Click on each workflow displayed and cancel it by clicking the **Cancel** button.

2. Rerun the rollback procedure from the start of section *Rollback ENM on Cloud* in *ENM on Cloud Upgrade Instructions (2/153 72-AOM 901 151)*.

Results

A successful rollback upon retry.

20.5.5 ENM Upgrade Workflow Hanging

> This section describes steps to take when an ENM upgrade workflow hangs and the VNF-LCM server log indicates that this is because of one or more servers failing to shut down.

ENM Upgrade Workflow Hanging

ENM Upgrade workflow is hanging and the VNF-LCM server log indicates that this is because of one or more servers failing to shut down.

1. If an **Upgrade Workflow** hangs, check the log file `/ericsson/3pp/jboss/standalone/log/server.log` on the `vnflaf-services` VM and look for the following message: `Still waiting for the shutdown of Pet VMs`.
2. Log on to the vCenter client URL `https://<vcenter_ip_vio_mgt>/ui/` as administrator `administrator@vsphere.local` with password `<vcenter_sso_password>`.
3. In the **Navigator** pane, click **Hosts and Clusters**.
4. Click and expand **vCenter Server**, then click and expand **Datacenter**, then select **Cluster**.
5. Select **Monitor** in the right panel, then select **Tasks & Events** and, in the left panel, select **Tasks**.
6. Investigate to see if any of the tasks have the following status message: `The operation is not allowed in the current state`.
7. Select that task related to the message if found.
8. Check the `Error stack` message in the bottom half of that panel for the following message: `The operation cannot be performed because VM migration is in progress`.
9. Identify the server instance name associated with the task above in **Navigator**.
10. Shut down the server instance found in step 9 through OpenStack.
    a. Log on to the VMS on the management network `<vms_ip_vio_mgt>` as user `root` with password `<vms_root_password>`.
    b. Source the `project.rcsource /vol1/senm/etc/<vio_os_project_name>_project.rc file`
    c. Shut down the server instance identified in step 9.

    ```
    openstack sever stop <server_instance_name>
    ```
11. Return to the VNF-LCM UI (using the below URL in the browser).

    ```
    [root@vnflaf-services ~http:///<external_ipv4_for_services_vm>/index.html#workfl →
    ows
    ```

Note: Replace the value for `<external_ipv4_for_services_vm>` in the URL with the IP address obtained earlier. If authentication is requested, log in using the username `vnfuser` and password `passw0rd` (with a zero).

12. Confirm that the **ENM Upgrade** workflow continues as normal.

20.6 Troubleshoot Problems with ENM Management

> This section describes how to troubleshoot and resolve problems related to the `manage_enm` script.

The described errors can occur with any of the `manage_enm` options - `start`,`stop` or `recover`.

20.6.1 Keystone Credentials Missing

> Troubleshoot a failure of the `manage_enm` script due to missing keystone credentials.

Cause

OpenStack RC file not sourced or not provided as an optional parameter.

The `manage_enm` script fails with an error similar to the example.

```
[root@vms ~]# manage_enm stop
INFO: Logging to /root/manage_enm.log
INFO: Recover ENM started
INFO: Authenticating with OpenStack
ManageEnmError - Failed loading Keystone credentials: Missing values for OS_AUTH_URL, OS_USERNAME, OS_PASSWORD, OS_IDENTITY_API_VERSION, OS_TENANT_ID, OS_ENDPOINT_TYPE.
[root@vms ~]#
```

Solution

1. Source the OpenStack RC file or provide the path to the file as a parameter to the `manage_enm` script.

   Example

   ```
   [root@vms ~]# manage_enm stop --rcfile /root/vio_123_project.rc
   ```

20.6.2 One or More Keystone Credentials Missing

Troubleshoot a failure of the `manage_enm` script because of a keystone error.

Cause

The `manage_enm` script fails with an error similar to the example.

```
[root@vms ~]# manage_enm start --rcfile /vol1/senm/etc/vio_5592_project.rc
INFO: Logging to /root/manage_enm.log
INFO: Recover ENM started
INFO: Authenticating with OpenStackManageEnmError - Failed loading Keystone cred →
entials: Missing values for OS_PROJECT_ID.
[root@vms ~]
```

Solution

1. Make sure that the correct OpenStack RC file is used and check that the missing credentials are added to the OpenStack RC file.

20.6.3 Manage ENM Script Fails to Authenticate with OpenStack

Troubleshoot a failure of the `manage_enm` script to authenticate with OpenStack.

Cause

The `manage_enm` script fails with an error similar to the example.

```
[root@vms ~]# manage_enm recover
INFO: Logging to /root/manage_enm.log
INFO: Recover ENM started
INFO: Authenticating with OpenStack
OpenStackError - Authentication with Openstack failed. Please check OpenStack se →
rvice is available.
[root@vms ~]#
```

Solution

1. Recover the OpenStack services by following the steps described in Restore OpenStack Services on page 10.

20.6.4 VNF-LCM Not Found

> Troubleshoot a failure of the `manage_enm` script to find VNF-LCM.

Cause

The hostname or IP passed to the `manage_enm` script in the optional parameter `--lcm` does not exist.

The `manage_enm` script fails with an error similar to the example.

```
[root@vms ~]# manage_enm recover --lcm 10.10.22.22
INFO: Logging to /root/manage_enm.log
usage: manage_enm.py recover [-h] [--rcfile RCFILE] [--lcm HOST] [--reason REASO →
N]
manage_enm.py recover: error: argument --lcm: Failed lookup of VNF-LCM host "10. →
10.22.22". Please enter a valid VNF-LCM host name or IP.
[root@vms ~]#
```

Solution

1. Make sure a valid VNF-LCM hostname or IP is provided.

20.6.5 Manage ENM Script Fails with 'ConnectionError'

> The following section troubleshoots connection errors with the `manage_enm` script.

The `manage_enm` script fails because of a connection error like in the following example.

```
INFO: Monitoring workflow instance ShutdownENM_20190206_100741
.....ConnectionError - HTTPConnectionPool(host='vnflaf-services', port=80): Max →
retries exceeded with url: /wfs/rest/progresssummaries/0b99cb94-29f7-11e9-9f71-f →
a163e760ba7 (Caused by
NewConnectionError('<requests.packages.urllib3.connection.HTTPConnection object →
at 0x7f26a94ac090>: Failed to establish a new connection: [Errno 113] No route t →
o host',))
```

1. Open the VNF-LCM UI in the browser using the following URL.

   ```
   http://<external_ip_for_services_vm>/index.html#workflows
   ```

   Note: Replace the value for `<external_ip_for_services_vm>` with the value corresponding to either the `<external_ipv4_for_services_vm>` or `<external_ipv6_for_services_vm>` parameter in the VNF-LCM SED.

2. Click the **Manage ENM - Shutdown** workflow under **Workflows** (the latest workflows are displayed under **Instance Activity**).
3. Confirm that the most recent **Manage ENM - Shutdown** workflow is complete.
4. Click the most recent workflow if it is not complete to view its progress and use the **Refresh** button in the top right-hand corner to track its progress until complete.
5. Rerun the `manage_enm` script with the stop action to ensure that the VNF-LCM VMs have been shut down.

20.7 Troubleshoot Problems with the vCenter VMware Integrated OpenStack Plugin

This section details how to troubleshoot issues with the vCenter VMware Integrated OpenStack plugin.

20.7.1 VIO Deployment Not Visible in vCenter GUI

This section describes how to recover if the VMware Integrated OpenStack (VIO) plugin is not visible in vCenter GUI.

Prerequisites

— VIO installed

— VIO plugin not visible in vCenter GUI.

Steps

1. Log out of vCenter client and log in again if the **VMware Integrated OpenStack** plugin is not visible in the **Menu** icon at the top of the vSphere client under **Inventories**.
2. Log out of vCenter client and log in again if the **VMware Integrated OpenStack** plugin is visible but OpenStack status is not **Running** under **VMware Integrated OpenStack** > **Deployment List** > **OpenStack Deployments**.
3. Follow section Restore OpenStack Services on page 10 if the **VMware Integrated OpenStack** plugin is still not visible.

20.7.2 Resolve OpenStack State Inconsistency

> This section describes how to fix an inconsistency between the OpenStack state reported by the VMware Integrated OpenStack plugin in vCenter and `viocli` on OMS.

This inconsistency is expressed by vCenter reporting that OpenStack is in an `Unavailable` (stopped) state while `viocli` reports OpenStack services as being in a `SUCCESS` (running) state. The solution to resolve this inconsistency is to reboot the OMS VM.

Prerequisites

— ENM SED document is available.

— VMS and OMS VMs are available and in a healthy state.

— A discrepancy in the reported OpenStack state between vCenter and OMS exists.

Steps

1. Log on to VMS on the management network `<vms_ip_vio_mgt>` as `root` user with the password `<vms_root_password>`.

   Note: Where `<vms_ip_vio_mgt>` and `<vms_root_password>` are parameters in the ENM SED.

2. Log on to the OMS.

   ```
   # ssh viouser@oms
   ```

3. Change to `root` user.

   ```
   $ sudo -i
   ```

4. Reboot the OMS VM.

   ```
   # reboot
   ```

   Note: After rebooting the VM, the session on OMS is disconnected.

5. Wait for the OMS VM to come back up, log on to it and check if the OpenStack state corresponds to the one reported in vCenter.

   ```
   $ sudo viocli deployment status --period 60
   ```

Note: Keep checking the status of the OpenStack services until all them are in a SUCCESS state.

20.8 Resolve File System Errors

This section describes the steps to resolve file system errors for a VM that may be encountered when the VM was restarted unexpectedly.

Prerequisites

The VM tries to boot and fails with file system errors.

Steps

1. Determine the file system name from the error seen in console of the VM.
2. Run the following command for each file system to repair the file systems for the VM.

   ```
   # fsck -f <device name> -y
   ```

   Note: Where <device name> is the name of the file system of the VM.
3. Reboot the VM.

Results

The VM is successfully recovered from the file system errors and it is running.

20.9 OMS Backup Script Failed ENM Alarm

This section describes the steps to take when an OMS Backup Script Failed ENM Alarm is raised.

Prerequisites

OMS Backup Script Failed alarm raised in the ENM Alarm Monitor.

Steps

1. Verify that connectivity to the OMS can be established.
   a. Log on to the VMS on the management network <vms_ip_vio_mgt> as user root with password <vms_root_password>.

b. Log on to the OMS.

```
[root@vms ~]# ssh viouser@oms
```

c. Follow the remaining steps to gracefully shut down and power on the OMS if connectivity cannot be established.

2. Verify that the disk is not full on the VMS.

a. Log on to the VMS on the management network <vms_ip_vio_mgt> as user root with password <vms_root_password>.

b. Run the following command to check if the disk is full.

```
[root@vms ~]#df -h
Filesystem                      Size  Used Avail Use% Mounted on
/dev/mapper/rhel-root            95G  3.8G   91G   4% /
devtmpfs                        1.9G     0  1.9G   0% /dev
tmpfs                           1.9G     0  1.9G   0% /dev/shm
tmpfs                           1.9G   20M  1.9G   2% /run
tmpfs                           1.9G     0  1.9G   0% /sys/fs/cgroup
/dev/sdb1                        99G    0G 1001G  79% /vol1
/dev/sda1                       497M  126M  372M  26% /boot
vENM_backup_env1                 95G  3.8G   91G   4% /BACKUP
tmpfs                           378M     0  378M   0% /run/user/0
```

Note: In the previous example, the disk space available on /vol1 is 0%. Disk space must be made available for an OMS backup to run successfully. FollowVMS Artifact Cleanup Procedure on page 102 and proceed with the next step.

3. Verify that no active viocli processes are running on the OMS.

a. Log on to the OMS.

```
[root@vms ~]# ssh viouser@oms
```

b. Run the following command to check for existing viocli processes.

```
[viouser@oms]# ps -ef | grep viocli
root 437 436 0 04:30 ? 00:00:00 sudo /usr/bin/viocli backup mgmt_se →
rver 10.0.0.5:/vol1/oms
root 438 437 0 04:30 ? 00:00:20 /usr/bin/python2.7 /usr/bin/viocli  →
backup mgmt_server 10.0.0.5:/vol1/oms
root 3827 3789 0 16:41 pts/0 00:00:00 grep --color=auto viocli
root 4318 4317 0 Sep02 ? 00:00:00 sudo /usr/bin/viocli backup mgmt_  →
server 10.0.0.5:/vol1/oms
root 4319 4318 0 Sep02 ? 00:03:35 /usr/bin/python2.7 /usr/bin/viocl →
i backup mgmt_server 10.0.0.5:/vol1/oms
```

Note: In the previous example, there are several viocli processes running. To completely remove the viocli the OMS must be gracefully shut down and powered back on. If there are viocli processes running on the OMS, follow the next step, otherwise skip the next step.

4. Gracefully shut down and power back on the OMS.

a. Log on to the vCenter client URL `https://<vcenter_ip_vio_mgt>/ui/` as administrator `administrator@vsphere.local` and password `<vcenter_sso_password>`.

b. Click **Hosts and Clusters** in the **Navigator** pane and click and expand the following items;**vCenter Server**, **Datacenter**, **Cluster**, and finally the vApp, **VIO**.

c. Right-click the virtual machine named **management-server** and under **Power** select **Shut Down Guest OS**.

d. Monitor the **management-server** by selecting the **Summary** tab and waiting for `Powered Off` to display in the top left of the main pane.

e. Right-click the VM once the VM is powered off and under **Power** select **Power On**.

f. Monitor the VM as detailed previously and wait for `Powered On` to display.

5. Verify that a successful backup can be taken.

a. Log on to the VMS on the management network `<vms_ip_vio_mgt>` as user root with password `<vms_root_password>`.

b. Run the following command.

```
[root@vms ~]# /opt/ericsson/senm/bin/oms_bur.sh -a backup -p {{ vms
_ip }}:/vol1/oms -r 7 -y || /opt/ericsson/senm/bin/send_FM_event.sh
 --managed_object_instance 'OMS (OpenStack Management Server)' --re
cord_type ALARM --event_type 'VIO ALARM' --perceived_severity CRITI
CAL --probable_cause 'VMS volume disk space full' --specific_proble
m 'OMS Backup Script Failed'"
```

c. Contact Ericsson Customer Support if a failure is shown in the output.

d. Verify that a backup was created by listing the contents of `/vol1/oms` and verify if the new backup files exist, named `vio_ms_<date>` and `vio_os_db_<date>` where `<date>` is today's date.

6. Resolve the alarm in the ENM Alarm Monitor provided the previous steps were successful .

a. Log on to ENM using the FQDN `httpd_fqdn` from the SED.

b. Under the section **Monitoring**, select **Alarm Monitor**.

c. Select **Add Topology Data** under **Network**.

d. Select **Search** and type in **`ManagementSystem`** in the new window that pops up.

e. Press **Enter** to search and select `ENM` from the results.

f. Click **Add** and then on the left-hand side click **Apply**.

Note: This lists all the alarms.

g. Find and select the alarm with the **OMS (OpenStack Management Server)** as the **Alarming Object** and click **Clear** and confirm.

20.10 Handle vSAN Health Alarm 'MTU Check (Ping with Large Packet Size)'

> This section describes how to handle a failed vSAN health check that raises alarm `vSAN health alarm 'MTU check (ping with large packet size)`.

The vSAN large ping health check sends jumbo frames from each ESXi host over the vSAN VMkernel interface `vmk1` to its peers on the other two ESXi hosts.

The health check runs even though Small Integrated ENM Multi-Technology does not use jumbo frames for vSAN traffic.

The jumbo frames are fragmented to handle the configured MTU size of 1500. If the alarm displays in vCenter, it can indicate a bad or badly seated uplink cable from one of the ESXi hosts to a physical switch.

The steps show the commands to use to identify the ESXi host and the uplink currently carrying vSAN traffic.

Note: Keep in mind a bad uplink cable may not always be the cause of the alarm.

Prerequisites

vSAN health check alarms appear in vCenter for `MTU check (ping with large packet size)`

Steps

1. Log on to the first ESXi host `<esxi_host1_vio_mgt_hostname>` as user `root` with password `<esxi_host1_mgt_password>` from the SED.
2. Confirm all uplinks used by the deployment have stable link status `Up`.

```
# esxcli network nic list
```

Note: Repeat the command at intervals to make sure that link status is stable.

3. Check if the vSAN interface can send jumbo frames to the vSAN interfaces on the other two ESXi hosts.

```
# vmkping -I vmk1 -s 9000 <esxi_host2_ip_vsan>
# vmkping -I vmk1 -s 9000 <esxi_host3_ip_vsan>
```

4. Repeat for the second ESXi host `<esxi_host2_vio_mgt_hostname>` and the third ESXi host `<esxi_host3_vio_mgt_hostname>`

   Note:
   - The ESXi host with a bad uplink assigned to vSAN traffic will not respond to the large pings.
   - The uplink currently carrying vSAN traffic can be identified with the `esxtop` command when logged on to the ESXi host with the issue.

     ```
     # esxtop
     ```

   - Press 'n' in `esxtop` to display the network screen. The `vmnic` assigned to `vmk1` is the uplink carrying vSAN traffic. Depending on the type of cable issue, the vSAN traffic may already have failed over to a good uplink so this uplink may be healthy.

5. Confirm all 4 10Gb uplinks for the host are pushed home to the host and physical switches.

   Note:
   - If the issue persists the uplink carrying vSAN traffic may be damaged.
   - Individual uplinks may be removed from the VDS in vCenter to help identify a bad uplink.
   - vSAN traffic fails over to another uplink.
   - Care is needed not to leave an ESXi host without at least one good uplink (preferably two good uplinks to separate switches) at all times. Without a working uplink, the host becomes isolated from the cluster and trigger a HA failover of all VMs on the host to the other two hosts.

Results

- ESXi host uplink cables were checked to make sure that none were badly seated.
- Large ping tests were run to help identify the cause of the vCenter alarms.

20.11 Handle vSAN Health Alarm 'vCenter State Is Authoritative'

> This section describes how to handle a failed vSAN health check that raises alarm `vCenter State is authoritative`.

This check verifies that all hosts in the vSAN cluster are using the current managing vCenter server as the source of truth for the cluster configuration, including the vSAN cluster membership list.

Prerequisites

— vSAN health check alarms display in vCenter for `vCenter State is authoritative`.

— No ESXi host is in maintenance mode.

Steps

1. Log on to the vCenter client URL `https://<vcenter_ip_vio_mgt>/ui/` as administrator `administrator@vsphere.local` with password `<vcenter_sso_password>`.
2. Click **Hosts and Clusters** in the **Navigator** pane.
3. Click and expand the following items **vCenter Server > Datacenter > Cluster > Monitor > vSAN > Skyline Health**.
4. Click **vCenter state is authoritative** and click **Update ESXi configuration**.

20.12 vSphere Web Client Slow to Launch

> This section describes the steps to perform if the vCenter client is slow to launch or respond.

Prerequisites

The vCenter client is slow to launch or respond.

Steps

1. Log on to the VMS on the management network `<vms_ip_vio_mgt>` as user `root` and password `<vms_root_password>`.

2. Connect through SSH to the vCSA appliance from VMS on management IP `<vcenter_ip_vio_mgt>` as user `root` and password `<vcenter_sso_password>` and restart the `vsphere-ui` and `vsphere-client` services.

```
Command> service-control --stop vsphere-ui
Command> service-control --start vsphere-ui
Command> service-control --stop vsphere-client
Command> service-control --start vsphere-client
Command> exit
```

Results

vCenter client responds as normal.

20.13 Known Issue: Service Unavailable

> This section describes the steps to recover from a Service Unavailable (HTTP 503) issue with workflows for Small Integrated ENM.

Prerequisites

Workflow indicates error `Service Unavailable (HTTP 503)`.

Steps

1. Log on to the VMS as `root` user.
2. Source the OpenStack project RC file.
3. List the volumes associated with the server in the stack that failed to update or delete.

   Note: The server name can be retrieved from the workflow log.

```
# openstack volume list |grep <server name>
```

4. Detach the volume from the server.

```
# openstack server remove volume <server name> <volume ID>
```

Note:
- Where `<volume ID>` is the volume ID from Step 3.
- Repeat for each volume, if there is more than one volume.

Results

The workflow can proceed.

20.14 Autostop Not Powering Off All VMs on Small Integrated ENM Transport Only

> This section describes a workaround for the issue of VMs still reporting as running after `autostop` is run on Small Integrated ENM Transport Only.

Prerequisites

VM(s) reported as running after running the `autostop` command.

Steps

1. Log on to the ESXi web client at `https://<esxi_host1_ip_vio_mgt>/` as user `root` with password `<esxi_host1_mgt_password>` from the SED.
2. In the left pane, click **Virtual Machines** and find the VMs reported as powered on by the GUI.
3. Check if the powered on VM is running by clicking the VM and checking its CPU and memory usage.
    a. If resources usage is not 0, gracefully shut down by clicking **Shut down** in the top pane and wait a few minutes for the VM to shut down.

        Note: If, after 15 minutes, the VM is still powered on contact Ericsson Support.

    b. If resources usage is 0, update the power state of the VM by clicking **Power off** in the top pane.
4. Repeat step 3 for each VM reported as running.

20.15 vSAN Object Health

> This procedure provides steps to resolve vSAN Object Health issue.

Prerequisites

vSAN object health issue is reported

Steps

1. Log on to the vCenter client URL `https://<vcenter_ip_vio_mgt>/ui/` as administrator `administrator@vsphere.local` with password `vcenter_sso_password`.
2. Go to **Navigator** and select the cluster inventory object **<vcenter_object_prefix>_CLUS**.
3. Select the **Monitor** tab and click **vSAN**.
4. Select **Skyline Health**.
5. Select **vSAN object health**.
6. Click the **Repair Objects Immediately** button and then click the **Retest** button to refresh the test result.
    a. Select **Resyncing Components** and confirm that there are no resyncing components.
    b. Select **Virtual Objects** and confirm vSAN Object Health is healthy for all VMs.
    c. Select **Physical Disks** and confirm that all disks are healthy.

Results

vSAN Object Health issue resolved.

Reference List

[1] *ENM on Cloud Deployment Instructions*, 7/1531-AOM 901 151
[2] *ENM on Cloud Upgrade Instructions*, 2/153 72-AOM 901 151
[3] *Small Integrated ENM Backup and Restore System Administrator Guide*, 2/1543-CNA 403 3456
[4] *ENM System Administrator Guide*, 1/1543-AOM 901 151
[5] *Small Integrated ENM Transport Only Installation Instructions*, 1/1531-CNA 403 3456
[6] *ENM Configuration System Administrator Guide*, 1/1543-AOM 901 151-1
[7] *ENM Troubleshooting Guide*, 1/159 01-AOM 901 151
[8] *Site Engineering Data for ENM on Cloud*, 2/1057-AOM 901 151
[9] *ENM Library Typographic Conventions*, 3/1551-FCK 101 05