HPE NFV Director

High Availability Solution Guide

Release 4.2

First Edition

Hewlett Packard
Enterprise

Notices

Legal notice

Confidential computer software. Valid license from HPE required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

The information contained herein is subject to change without notice. The only warranties for HPE products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HPE shall not be liable for technical or editorial errors or omissions contained herein.

Printed in the US

Trademarks

Linux is the registered trademark of Linus Torvalds in the U.S. and other countries.
Oracle and Java are registered trademarks of Oracle and/or its affiliates.

Adobe®, Acrobat® and PostScript® are trademarks of Adobe Systems Incorporated.

HP-UX Release 10.20 and later and HP-UX Release 11.00 and later (in both 32 and 64-bit configurations) on all HP 9000 computers are Open Group UNIX 95 branded products.

Microsoft®, Internet Explorer, Windows®, Windows Server 2007®, Windows XP®, and Windows 7® are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

Firefox® is a registered trademark of the Mozilla Foundation.

Google Chrome® is a trademark of Google Inc.

Oracle® is a registered U.S. trademark of Oracle Corporation, Redwood City, California.

EnterpriseDB® is a registered trademark of EnterpriseDB.

Postgres Plus® Advanced Server is a registered U.S. trademark of EnterpriseDB.

UNIX® is a registered trademark of The Open Group.

X/Open® is a registered trademark, and the X device is a trademark of X/Open Company Ltd. in the UK and other countries.

Red Hat® is a registered trademark of the Red Hat Company.

Apache CouchDB, CouchDB, and the project logo are trademarks of The Apache Software Foundation

Node.js project. Joyent® and Joyent's logo are registered trademarks of Joyent, Inc

VMware ESX, VMWare ESXi, VMWare vCenter and VMWare vSphere are either registered trademarks or trademarks of VMware, Inc. in the United States and other jurisdictions

Contents

List of tables

List of figures

Preface

About this guide

This guide NFVD HA Installation provides all the needed information to have HP NFV Director set up and maintained in a High Availability configuration.

Audience

This guide is intended for any stakeholder that needs to install and configure NFVD for a production environment in High Available mode. It is recommended that the person is knowledgeable in basic Linux and Oracle administration to use this document.

Document history

Table 1: Document history

Edition	Date	Description
1.0	18 January 2017	First version

References

Table 2: References

Document	Description
NFV Director Installation and Configuration Guide	Installation of a simplex (non-HA) instance of NFVD
NFV Director Administration Guide	Administration of a simplex (non-HA) instance of NFVD
NFV Director Troubleshooting Guide	Troubleshooting of a simplex (non-HA) instance of NFVD
SIS_11.31_SiteScopeDeployment	SiteScope Deployment Manual
SIS_11_31_Failover	Installation of a SiteScope hot standby server

Chapter 1
NFV Director HA Architecture

The High Availability (HA) architecture requires two copies of each NFVD SW process running on separate VM so that on the eventual failure of one process its mirror process takes over.

Ideally all process should be "active/active" with a load balancer in front or a virtual IP (VIP) so that each process is not only protected from an eventual failure but the load is distributed.

However, not all of the current versions of the underlying SW allow such a configuration; as a result there are some processes that run in "active/passive" mode.

There are two types of NFVD processes:

- Active / Active : N processes are active at the same time and a load balancer in front distributes the load (either with sticky session or round robin mechanism)
- Active / Passive: Only one primary process is active and there is a standby process that is activated in the event of the failure of the primary

The following figure illustrates how the various NFVD processes are set up in an HA configuration.

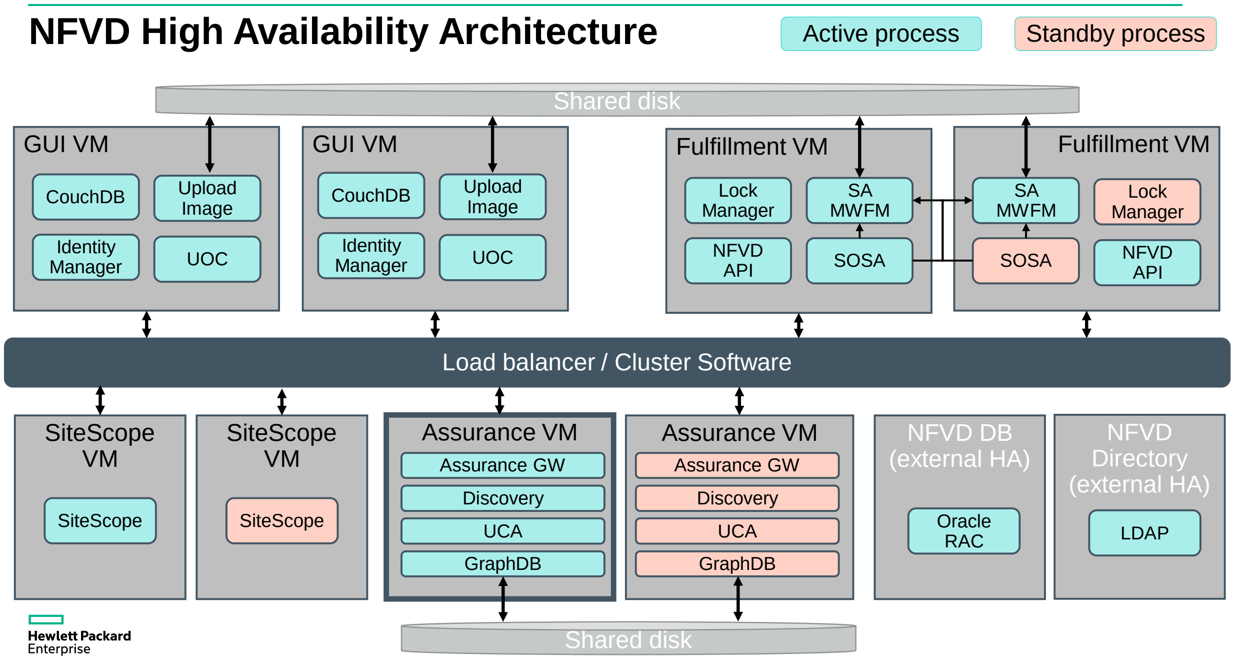

Figure 1: NFVD High Availability Configuration

In this document, it is assumed that the hosts running the components illustrated above are named as follows:

- GUI VMs: gui-node1 & gui-node2
- Fulfillment VMs: ff-node1 & ff-node2
- SiteScope VMs:sis-node1 & sis-node2
- Assurance VMs: aa-node1 & aa-node2

Chapter 2
NFVD Director HA Installation

2.1 Overview

Installing NFVD in an HA environment involves the following steps:

1. Reviewing the NFVD HA Architecture.
    - Please refer to Chapter 1 "NFV Director HA Architecture" to become familiar with VMs and the naming convention used in the HA setup

2. Installing and configuring the prerequisite software for NFVD HA installation.
    - Please refer to Section 2.3 "Prerequisites"

3. Configuring shared disks for image repository and UCA EBC
    - Please refer to Section 2.3.7 "Configuring shared disks"

4. Installing fully functional "primary" and "secondary" NFVD platform (Fulfillment VM, Assurance VM, GUI VM)
    - Please refer to Section 2.4 "Primary and Secondary Platform Preparation"

5. Connecting the secondary platform with the primary platform (using Virtual IPs when required)
    - Please refer to Section 2.5 "Connection of the Primary and Secondary Platforms"

6. NFVD HA Cluster configuration.
    - Please refer to Section 2.6 "Heartbeat Cluster Configuration"

7. Making the NFVD HA platform operational.
    - Please refer to Section 2.7 "Starting NFVD HA Platform"

8. Connectivity and Load Balancer Configuration
    - Please refer to Section 2.8 "Load Balancer Configurations"

9. Configuring Monitoring tools.
    - Please refer to Section 2.9 "Configuring Monitoring Tools"

10. NFVD HA Setup Validation.
    - Please refer to Section 2.10 "Validation of the NFVD HA setup"

11. Discover VIM resources.
    - Please refer to Chapter 3 "Integrating with the VIM to discover resources."

The setup is now ready for VNF deployments.

2.2 NFVD Director HA example rpm

The various scripts used for the NFV Director HA configuration are included as an example in the ha-example rpm. The contents are as follows:

```
/opt/HPE/nfvd/solutions/ha-example
/opt/HPE/nfvd/solutions/ha-example/aa check.sh
/opt/HPE/nfvd/solutions/ha-example/aa_maintenance_mode.sh
/opt/HPE/nfvd/solutions/ha-example/gui maintenance mode.sh
/opt/HPE/nfvd/solutions/ha-example/sis check.sh
/opt/HPE/nfvd/solutions/ha-example/sis_check_primary.sh
/opt/HPE/nfvd/solutions/ha-example/sischkprm
```

2.3 Prerequisites

2.3.1 Infrastructure Requirements

In order to guarantee High Availability at a system leveler, it is assumed that any HA deployment of NFV-D will be on a highly available infrastructure. In particular:

- The hardware infrastructure supporting NFV-D is fully redundant, with dual power supplies, ventilation, etc. at a rack level, and with primary/secondary servers isolated at a data center level
- Physical compute nodes hosting virtual servers are set up in a high availability cluster, so that a failure of a compute node will trigger the transfer of any hosted VMs on this node to be re-located on a new compute node
- Network connections between NFV-D components and external components, as well as connections internal to NFV-D, are redundant. This is especially important for networks hosting "heartbeat" traffic between nodes, in order to avoid "split brain" occurrences leading to data corruption.
- Customer-provided services, such as database or LDAP/AD directory services, are set up in a high-availability mode, transparent to the NFV-D applications.

2.3.2 Connectivity to External Components

Each server (primary and secondary) in the NFVD will require connectivity to hosts where the following customer-supplied components are deployed:

- Oracle database component (Oracle 12c) in order to deploy the NFVD data model and store persistent data. Oracle RAC is the recommended HA configuration.

- LDAPv3 directory server implementation. Typical examples are:
  - OpenLDAP without SSL connection.
  - ActiveDirectory with SSL connection.

  The recommended HA configuration for the directory server is multi-master

- An external load balancer (such as F5 or HAproxy). The external load balancer should also be set up in an HA configuration.

From NFVD standpoint, there is no constraint on how these components are actually deployed, either through physical or virtual hosts, either collocated or not collocated, as long as they meet connectivity requirements.

Notwithstanding, in order to have a complete HA solution, it is recommended that following components are deployed with following topology (but this is not a prerequisite as such to install NFVD HA):

- Oracle RAC
- LDAP or ActiveDirectory multi-master
- F5 or HAproxy High Availability

2.3.3 Allocation of Virtual IPs

The NFVD system requires physical IP addresses for each of the servers on which it will be deployed. IN addition, several Virtual IP Addresses are required for the components illustrated in Figure 1:

- GUI
- NFVD API
- SiteScope
- LockManager/SOSA
- Assurance Gateway

The table below describes the HA mechanism used for each of these components:

Table 3: NFVD Component Virtual IP providers

NFVD Component	Provider
GUI	External Load Balancer such as F5 or HAProxy. The Virtual IP is configured in external Load Balancer component
NFVD API	External Load Balancer such as F5 or HAProxy. The Virtual IP is configured in external Load Balancer component
SiteScope	External Load Balancer such as F5 or HAProxy. The Virtual IP is configured in external Load Balancer component
SOSA / Lock Manager	Heartbeat Service or equivalent cluster software. The Virtual IP is configured in the Cluster Software component.
Assurance Gateway / UCA	Heartbeat Service or equivalent clustering software. The Virtual IP is configured in the Cluster Software component.

Of course there needs to be connectivity between these various VIPs.

2.3.4 Installation of Cluster Software

The Assurance Gateway component and the SOSA/Lock Manager component require cluster software to be installed on their respective servers in order to provide the Virtual IP and failover mechanism for these components.

The NFVD HA configuration requires that the Fulfillment VM and Assurance VM primary servers and their respective standby servers be clustered to provide High Availability for the services they provide.

In this document, as a reference implementation, the Linux heartbeat daemon is used as a cluster software to provide High Availability for these components. However, the customer can chose to use other clustering solutions that provide a Virtual IP and failover mechanism.

Heartbeat is a daemon that provides cluster infrastructure (communication and membership) services to its clients. This allows clients to know about the presence (or disappearance!) of peer processes on other machines and to easily exchange messages with them

2.3.4.1 Heartbeat Installation on Fulfillment VMs

Each of the Fulfillment VMs requires the heartbeat daemon to be installed:

```
[root@ff-node1 ~]# yum install heartbeat
[root@ff-node2 ~]# yum install heartbeat
```

Note: If the yum utility cannot install the heartbeat service then heartbeat software needs to be installed manually. The list of rpms to be installed is provided in Section Error! Reference source not found.

2.3.4.2 Heartbeat Installation on Assurance VMs

Each of the Assurance VMs requires the heartbeat daemon to be installed:

```
[root@aa-node1 ~]# yum install heartbeat
[root@aa-node2 ~]# yum install heartbeat
```

2.3.4.3 NFVD HA VMs in an OpenStack Environment

If the NFVD HA VMs are to be installed on virtual machines within an OpenStack environment, then the VIPs associated with the Fulfillment VM cluster and Assurance VM Cluster need to "allowed" IPs for the ports of the internal network of the associated servers.

As a reference, here are the required commands that should be done in OpenStack to configure this:

```
#Create a port: (Preferably using an IP from the discovery range)
neutron port-create HA-private  --fixed-ip ip address=<vip>

#Get the ports Id attached to each server
neutron port-list | grep <internal-ip-vm-1>
neutron port-list | grep <internal-ip-vm-2>

#Attach the new IP to the port on each server
neutron port-update   fc211741-7a10-4ab0-a13b-c1de28aba6df --allowed_address_pairs list=true
type=dict ip address=<vip>
neutron port-update  c12fefaa-cee1-4913-ade1-620d4b3acc99 --allowed address pairs list=true
type=dict ip_address=<vip>
```

It is also required to configure the VMs with fixed IP addresses when creating them in OpenStack, and to configure the IPs as fixed within the VM operating systems.

2.3.4.4 **Generating** Shared Keys for the Heartbeat Clusters

A shared key is required for each heartbeat cluster in order to allow the heartbeat daemons to communicate between cluster members

The following commands are used to generate the shared keys:

```
[root@ff-node ~]$ dd if=/dev/urandom bs=512 count=1 | openssl md5
1+0 records in
1+0 records out
512 bytes (512 B) copied, 0.000119239 s, 4.3 MB/s
(stdin)= c9f81b396a4e74278faf02d04c60f16f
```

Separate shared keys should be generated for the Fulfillment VM cluster and the Assurance VM cluster.

2.3.5 Installation of External Load Balancer Software.

The GUI, NFVD API and SiteScope components require a load balancer software to be installed (externally to NFVD VMs) to provide the Virtual IP and failover mechanism for these components.

In this document, as a reference implementation, the Linux HAProxy daemon is used as a load balancer software to provide High Availability for these components. However, the customer can chose to use other clustering solutions that provide a Virtual IP and failover mechanism.

HAProxy or High Availability Proxy, is a popular open source software TCP/HTTP Load Balancer and proxying solution which can be run on Linux, Solaris, and FreeBSD. Its most common use is to improve the performance and reliability of a server environment by distributing the workload across multiple servers.

```
[root@LB-node ~]# yum install haproxy
```

Information is also provided for using the F5 Local Traffic Manager as the load balancer software.

2.3.6 Installation of the xinetd service

The xinetd (extended Internet daemon) is an open source super-server daemon which is required to provide NFVD HA capabilities. It is used to register the health-monitoring scripts that are invoked by the external load balancer.

The xinetd service needs to be installed on all the NFVD HA VMs.

By default on RHEL the xinetd service is pre-installed. If it is not installed then the xinetd should be installed on all NFVD HA VMs using the commands described in the following sections.

2.3.6.1 Install xinetd software on GUI VMs

```
[root@gui-node1 ~]# yum install xinetd
[root@gui-node2 ~]# yum install xinetd
```

2.3.6.2 Install xinetd software on Fulfillment VMs

```
[root@ff-node1 ~]# yum install xinetd
[root@ff-node2 ~]# yum install xinetd
```

2.3.6.3 Install xinetd software on SiteScope VMs

```
[root@sis-node1 ~]# yum install xinetd
[root@sis-node2 ~]# yum install xinetd
```

2.3.6.4 Install xinetd software on Assurance VMs

```
[root@aa-node1 ~]# yum install xinetd
[root@aa-node2 ~]# yum install xinetd
```

2.3.7 Configuring shared disks

When NFVD is deployed in an HA configuration, certain data need to be shared between VMs:

- The GUI and Fulfillment VMs (on both primary and secondary hosts) need to share a directory in which all VNF VM images will be stored.
- The Assurance VMs on the primary and secondary nodes need to share the same UCA EBC data repository

There are three methods possible to provide shared storage for the NFVD components. These are listed in decreasing order of preference:

- Preferred solution:
  - An external storage system provides a single volume through NFS and all VMs that need to share the storage mount the same volume

- Next best solution:
  - An external storage system (or OpenStack cinder) provides 2 volumes; each VM mounts one of the volumes
  - Each VM configures glusterFS to replicate and sync data between the 2 volumes

- Last-chance option:
  - Each VM defines a volume using the local disk
  - Each VM configures glusterFS to replicate and sync data between the 2 volumes

For the remainder of this document, it is assumed that the (preferred) NFS option has been chosen.

2.3.7.1 NFS Configuration

The following packages need to be installed in order to be able to mount the NFS exported directories shared by NFS server. These packages should be installed on the GUI VM, the Fulfillment VM, and the Assurances VM:

```
[root@ff-node ~]# yum install nfs-utils rpcbind
```

Once installed the RPC service should be started:

```
[root@ff-node ~]# service rpcbind start
```

Create a folder in the filesystem and mount the NFS exported directory:

```
[root@ff-node ~]# mkdir -p /mnt/nfs
[root@ff-node ~]# mount <NFS_server_IP>:<NFS_folder_shared_by_NFS_server> /mnt/nfs
```

If, for example, the server is, 10.75.14.23, and it shares the '/opt/nfs' folder, the following should be executed:

```
[root@ff-node ~]# mount 10.75.14.23:/opt/nfs /mnt/nfs
```

In order for the VMs to mount the NFS exported directory on reboot, an entry similar to the following should be added to the /etc/fstab file:

```
10.75.14.23:/opt/nfs       /mnt/nfs     nfs    defaults            0 0
```

The hostnames (NFS server and clients) should be known on both sides and be fully qualified in /etc/hosts on each server

The proper Domain must be set in /etc/idmapd.conf file of the NFS server and clients. For example:

```
Domain = gre.hpecorp.net
```

2.3.7.2 Sharing the VNF Image Service Repository

The image directory is configured in an NFS mount point, shared between the GUI and Fulfillment VMs.

A typical configuration is to use /image_repository with a size of 100 GB

- Note: do not include /var/opt/uoc2/server/public in the path of NFS mount point

Special considerations need to be paid to the protections on the NFS shared folder, as both the Fulfillment user (root) on the Fulfillment VMs and the GUI user on the GUI VMs need read/write access to this folder.

A possible procedure to assigning the correct permissions is as follows.

1. In the NFS server and the GUI VMs, add a ‘uoc’ user:

```
[root@hpe-nfs-server ~]# adduser uoc
[root@gui-node1 ~]# adduser uoc
[root@gui-node2 ~]# adduser uoc
```

2. The uid (user ID) and gid (group ID) for the uoc user in the NFS server must be the same as in both GUI nodes. In the GUI nodes, the uid/gid for the uoc user can be found in /etc/passwd:

```
[root@gui-node1 ~]# cat /etc/passwd
[…]
uoc:x:50010:4012::/export/home/uoc:/bin/bash
```

3. On the NFS server, search the /etc/passwd file for the uoc line. Modify that line with the correct uid/gid obtained from the GUI nodes:

```
uoc:x:50010:4012::/home/uoc:/bin/bash
```

4. Change (recursively) folder/files permissions (in GUI nodes) for the NFS shared folder

```
[root@gui-node1 ~]# chown -R uoc.root /image_repository
[root@gui-node2 ~]# chown -R uoc.root /image_repository
```

Once the previous steps have been completed, both the Fulfillment user (root user) on the Fulfillment VMs and the GUI user (uoc user) on the GUI VMs should be able to write/read into NFS shared folder.

To validate, check that files can be created as uoc user from the GUI nodes and root user from the Fulfillment VM nodes:

```
[uoc@gui-node1 ~]$ touch /image_repository/test1
[uoc@gui-node2 ~]$ touch /image_repository/test2
[root@ff-node1 ~]$ touch /image_repository/test3
[root@ff-node2 ~]$ touch /image_repository/test4
```

2.3.7.3 Sharing the UCA-EBC data repository

On the Assurance VMs, the /var/opt/UCA-EBC directory must be in sync for both the VMs. The size of directory is typically 50 GB.

The following steps are required to set up the synchronization using NFS:

1. In NFS Server and AA nodes, add an ‘uca’ user:

```
[root@hpe-nfs-server ~]# adduser uca
[root@aa-node1 ~]# adduser uca
[root@aa-node2 ~]# adduser uca
```

2. The uid (user ID) and gid (group ID) for the 'uca' user needs to be the same in NFS server and both Assurance nodes.

3. Create a folder in the Assurance nodes filesystems and mount the NFS exported directory (assuming directory to export is also /var/opt/UCA-EBC on NFS server):

```
[root@aa-node1 ~]# mkdir -p /var/opt/UCA-EBC
[root@aa-node1 ~]# mount <NFS_server_IP>:/var/opt/UCA-EBC /var/opt/UCA-EBC
[root@aa-node1 ~]# chown -R uca:uca /var/opt/UCA-EBC

[root@aa-node2 ~]# mkdir -p /var/opt/UCA-EBC
[root@aa-node2 ~]# mount <NFS_server_IP>:/var/opt/UCA-EBC /var/opt/UCA-EBC
[root@aa-node2 ~]# chown -R uca:uca /var/opt/UCA-EBC
```

4. On the Assurance nodes, add the following entry in the /etc/fstab file:

```
<NFS_server_IP>:/var/opt/UCA-EBC  /      var/opt/UCA-EBC nfs defaults      0 0
```

5. On the Assurance nodes, as uca user, check that a file can be created in the shared data repository and that this file has uca:uca permissions:

```
[uca@aa-node1 ~]$ touch /var/opt/UCA-EBC/test1
[uca@aa-node2 ~]$ touch /var/opt/UCA-EBC/test2
```

2.4 Primary and Secondary Platform Preparation

This chapter describes the initial installation of the primary and secondary NFVD HA platforms.

2.4.1 Primary Platform Installation

Installation of the primary platform is done exactly the same way as described in the "*NFVD 4.2 Installation and Configuration Guide*".

```
# /opt/HPE/nfvd/install/nfvd-install.sh /kits/archives
```

Please follow the instructions in this guide (including "Post Installation" tasks and Special instructions (if any) in "*NFVD 4.2 Release Notes*".

Note:
When the installer asks for the "Image Uploader Service repository", specify the folder configured in Chapter 2.3.7, Section 2.3.7.2 "Sharing the VNF Image Service Repository" for example, /image_repository)

2.4.2 Secondary Platform Preparation

This section describes the preparation of a secondary NFVD platform (consisting of the secondary GUI, Fulfillment, Assurance, and SiteScope VMs). The installation requires the use of a specific mode of the NFVD installer so that the Secondary platform can later be connected to the Primary platform installed according to Section 2.4.1.

The installation is almost identical to that of the primary platform. Follow the instructions in the NFVD 4.2 Installation and Configuration Guide, but use the '-s' option of the nfvd-install.sh tool:

```
# /opt/HPE/nfvd/install/nfvd-install.sh -s /kits/archives
```

Note:

1. When the installer asks for the different hosts composing the secondary platform, use the actual IP addresses and hostnames of the secondary VMs, and not the Virtual IPs or Virtual Hostnames (these will be configured during the "connection" described in Chapter +++"Connecting NFVD platforms").

2. When the installer asks for the "Image Uploader Service repository", specify the folder configured in Chapter 3, Section 3.2 "Sharing the VNF Image Service Repository" for example, /image_repository)

3. At the end of the secondary platform installation, it is normal that no NFVD processes are running on the different hosts composing the secondary platform.

The Post Installation steps (section 3.1 to 3.4 of Chapter 3) described in the "*NFVD 4.2 Installation and Configuration Guide"* should also be executed. However, the sections "3.5 - Configuring NFVD domain user" and "3.6 – Verifying NFVD installation" have to be skipped.

2.5 Connection of the Primary and Secondary Platforms

This chapter describes the configuration steps necessary to connect the primary and secondary platforms. This step requires Virtual IPs to be supplied for the following components:

- LockManager/SOSA
- NFVD API
- Assurance Gateway
- SiteScope
- GUI

2.5.1 Reconfiguration of the Secondary Platform with Virtual IPs

On the <INSTALLER_HOST> of the **secondary** platform, login as "root" and execute the following command:

```
# /opt/HPE/nfvd/install/nfvd-reconfigure.sh -s
```

This script will first prompt to confirm the parameters specified during the initial secondary platform install. The default value will automatically be filled with the value that was set at installation time in Section 2.4.2 above. If there is no change, just hit "Enter".

Next, the script will prompt for the following Virtual IP addresses:

- LockManager/SOSA (when prompted for the LockManager VIP and the SOSA VIP, enter the same VIP)
- NFVD API
- Assurance Gateway
- SiteScope
- GUI

2.5.2 Reconfiguration of the Primary Platform with Virtual IPs

On the <INSTALLER_HOST> of the **primary** platform, login as "root" and execute the following command:

```
# /opt/HPE/nfvd/install/nfvd-reconfigure.sh -p
```

(i.e. use nfvd-reconfigure.sh with the -p option instead of the -s option)

Again, the script will prompt to confirm the parameters, and then prompt for the Virtual IP addresses.

Note : At the end of reconfiguration all the NFVD processes on all NFVD VMs will be down.

2.5.3 Configuration of SiteScope

Note: Ensure that while importing the SiteScope license, you import the SiteScope Failover license included with the SiteScope OSI Pack SW License.

Detailed instructions to configure the SiteScope Failover node can be found in the document "*SIS_11_31_Failover.pdf*", available on the HPE software support site https://softwaresupport.hpe.com/

A brief summary of the necessary configuration steps is provided in the following sections.

2.5.3.1 Restart SiteScope Primary and SiteScope Failover.

After Reconfiguration check the status of Primary and Failover SiteScope processes

- Primary SiteScope

```
[root@sis-node1 ~]# /opt/HPE/nfvd/bin/nfv-director.sh -c sitescope status
```

If it returns [Not Running], then start SiteScope

```
[root@sis-node1 ~]# /opt/HPE/nfvd/bin/nfv-director.sh -c sitescope start
```

The expected output is

```
SiteScope                              : [Started]
```

- Secondary SiteScope

```
[root@sis-node1 ~]# /opt/HPE/nfvd/bin/nfv-director.sh -c sitescope status
```

If it returns [Not Running], then start SiteScope

```
[root@sis-node1 ~]# /opt/HPE/nfvd/bin/nfv-director.sh -c sitescope start
```

The expected output is

```
SiteScope                              : [Started]
```

2.5.3.2 Creation of a Failover Profile

In the Failover (secondary) node UI, go to Preferences > High Availability Preferences.

In the right panel, click New Profile to open the New Failover Profile dialog.

Specify the settings as required [sample in screenshot below], and then click OK,

The value “Host” is the IP address of the Primary SiteScope.

The value “Port” is the port on which Sitescope is running i.e 18888

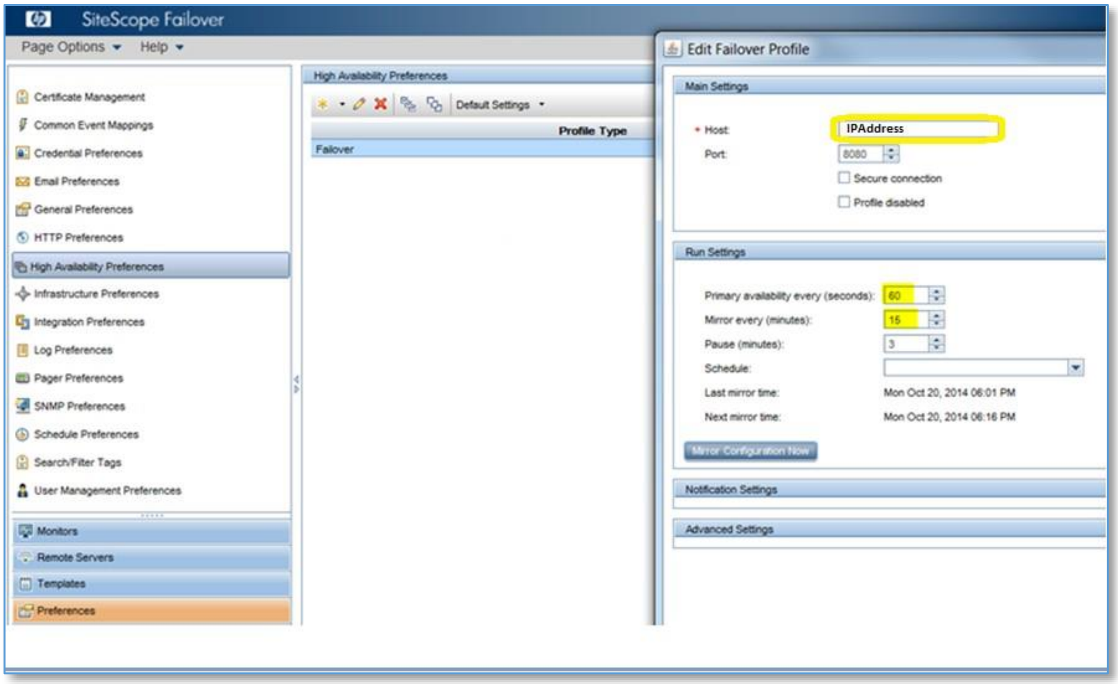

* Host: IPAddress
Port: 8080
□ Secure connection
□ Profile disabled

Primary availability every (seconds): 60
Mirror every (minutes): 15
Pause (minutes): 3
Schedule:
Last mirror time: Mon Oct 20, 2014 06:01 PM
Next mirror time: Mon Oct 20, 2014 06:16 PM

Figure 2: SiteScope Failover Profile Preferences

2.5.3.3 Verification of Failover node settings

Login to the Primary SiteScope UI using the Primary node IP.

Go to Preferences > High Availability Preferences. Select Default Settings > Test.

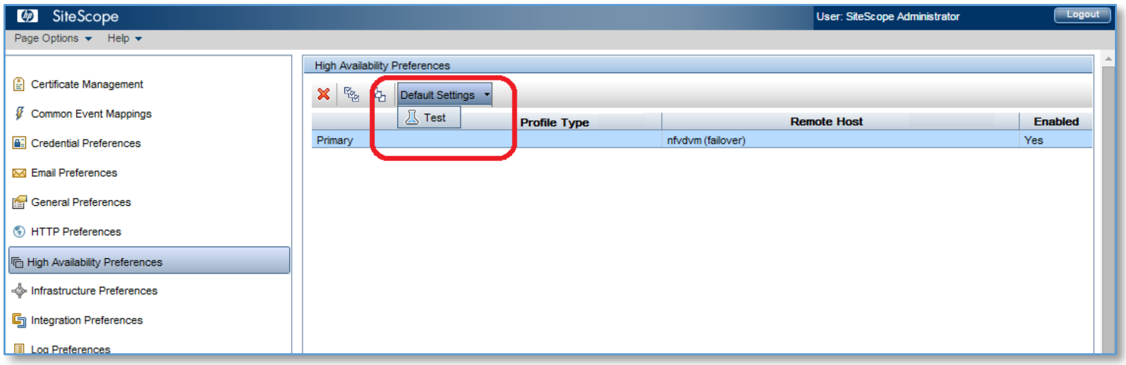

Figure 3: Testing SiteScope High Availability Setup

The test results should resemble the following:

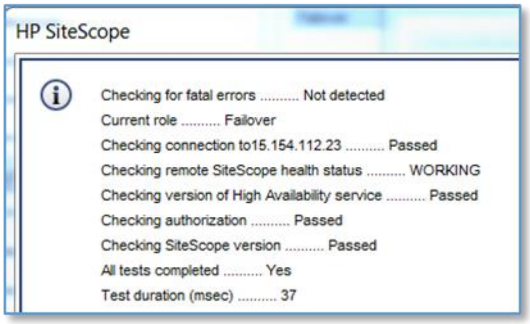
HP SiteScope

```
Checking for fatal errors .......... Not detected
Current role .......... Failover
Checking connection to15.154.112.23 .......... Passed
Checking remote SiteScope health status .......... WORKING
Checking version of High Availability service .......... Passed
Checking authorization .......... Passed
Checking SiteScope version .......... Passed
All tests completed .......... Yes
Test duration (msec) .......... 37
```

Figure 4: SiteScope High Availability Test Results

2.5.3.4 Configuration of the SiteScope SNMP Alarm target

When SiteScope monitors detect alarms, these are forwarded to the active Assurance VM with SNMP. On both the primary and failover SiteScope nodes, perform the following:

- Go to Preferences > SNMP Preferences
- Edit the SNMP Target entry and provide the Virtual IP of the Assurance Gateway component
- Click on OK

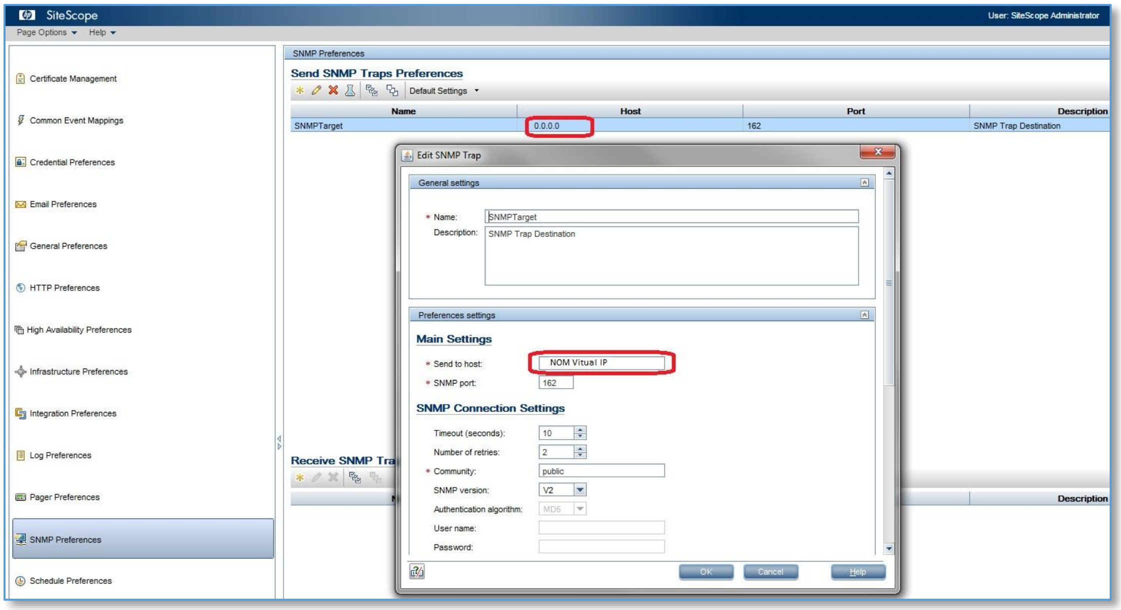

Send SNMP Traps Preferences

Name	Host	Port	Description
SNMPTarget	0.0.0.0	162	SNMP Trap Destination

Edit SNMP Trap

General settings

* Name: SNMPTarget
Description: SNMP Trap Destination

Preferences settings

Main Settings

* Send to host: NOM Vitual IP
* SNMP port: 162

SNMP Connection Settings

Timeout (seconds): 10
Number of retries: 2
* Community: public
SNMP version: V2
Authentication algorithm: MD5
User name:
Password:

Figure 5: Configuration of the SiteScope SNMP target

2.5.3.5 Configuration of SiteScope Receiver URL for Data Integration Preferences

On both the primary and failover SiteScope nodes, update Receiver URL value in Data Integration Preferences tab by replacing <Assurance VM1> IP address with <Assurance Heartbeat VIP> IP address.

- Click on Preferences > Integration Preferences.
- A record by name '<diname>' will be created. Open the same.
- In "Data Integration Preferences Settings" data, update Receiver URL value:

  http://<AssuranceHeartbeatVIP>:18080/nfvd/kpimetrics/

2.6 Heartbeat Cluster Configuration

This chapter describes how the Linux heartbeat daemon can be set up in the NFVD Fulfilment and Assurance components to provide High Availability for these components.

2.6.1 Fulfilment Heartbeat Configuration

A typical High Availability configuration for the NFVD Fulfilment component is shown in the figure below:

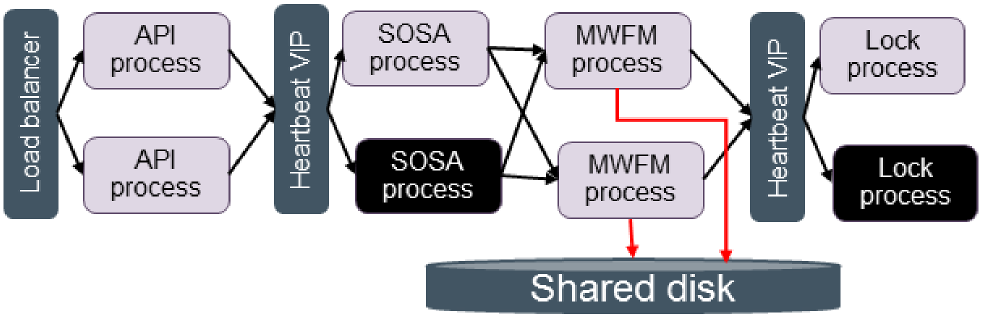

Figure 6: Fulfillment High Availability Architecture

- There are two active/active instances of the API process on the two Fulfilment VMs; this setup is described in Section 2.8.2
- The Linux heartbeat daemon is used to provide a Virtual IP to the SOSA and Lock processes. Only one instance of these processes can be operational at a time. Thus if the active SOSA or lock processes fail on one Fulfilment VM, they are started up on the other Fulfilment VM
- SOSA receives service requests from the FF API, and load-balances these requests between the two active HPW Service Activator MWFM instances
- The Service Activator MWFM instances are completely independent. If one of the instances fails, then workflows in progress on that instance are paused, and will resume at the current step when the MWFM instance recovers. IN the meantime, SOSA will direct new requests to the remaining MWFM instance.

2.6.1.1 Fulfilment Heartbeat Configuration Files

The following steps should be carried out on both Fulfilment VMs:

First copy the init script for SOSA and LockManager to the appropriate place:

```
cp /opt/HPE/nfvd/fulfillment/scripts/activator-ep /etc/init.d/
chmod +x /etc/init.d/activator-ep
```

Then configure the following files:

1. /etc/ha.d/authkeys

Set the mode of this file to "600"; its contents should be:

```
auth 2
2 sha1 <shared_key>
```

Where <shared_key> is the key generated with the tool as described in Section 2.3.4.4 "Generating Shared Keys for the Heartbeat Clusters"

2. /etc/ha.d/ha.cf

The content of this file should be this:

```
logfile /var/log/heartbeat.log
logfacility local0
keepalive 2
deadtime 30
initdead 120
ucast eth0 <other node in Fulfillment cluster>
udpport 694
auto_failback off
node <ff-node1>
node <ff-node2>
```

<ff-node1> and <ff-node2> should be the host names obtained by executing the command "uname –n" on the two fulfilment nodes.

The recommended configuration is to not allow the failback of resources when the primary node goes up again. This can be changed with using the auto_failback configuration to "on".

3. /etc/ha.d/haresources

```
<primary node in Fulfillment cluster> IPaddr::<vip>/24/eth0:0 activator-ep
```

The <vip> should be the Virtual IP for the Fulfilment heartbeat cluster.

2.6.2 Assurance Heartbeat Configuration

A typical High Availability configuration for the NFVD Assurance component is shown in the figure below:

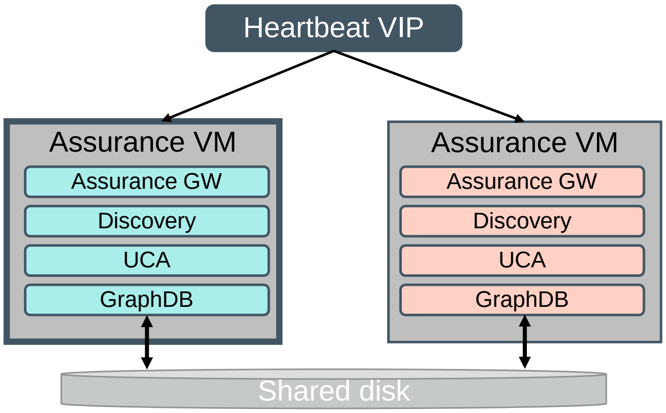

Figure 7: Assurance High Availability Architecture

The Linux heartbeat daemon is used to provide a Virtual IP to the Assurance Gateway process. The Assurance Gateway then directs traffic and requests to the correlation (UCA + GraphDB) and Discovery components. Only one instance of these processes can be operational at a time. Thus if any of these processes fail on the active Assurance VM, they are started up on the standby Assurance VM.

2.6.2.1 Assurance Heartbeat Configuration Files

The following steps should be carried out on both Assurance VMs:

First copy the init script for Assurance to the appropriate place:

```
cp /opt/HPE/nfvd/bin/nfvd /etc/init.d/
chmod +x /etc/init.d/nfvd
```

Then configure the following files:

1. /etc/ha.d/authkeys

Set the mode of this file to “600”; its contents should be:

```
auth 2
2 sha1 <shared_key>
```

Where <shared_key> is the key generated with the tool as described in Section 2.3.4.4 “Generating Shared Keys for the Heartbeat Clusters”

2. /etc/ha.d/ha.cf

The content of this file should be this:

```
logfile /var/log/heartbeat.log
logfacility local0
keepalive 2
deadtime 30
initdead 120
ucast eth0 <other node in Assurance cluster>
udpport 694
auto_failback off
node <aa-node1>
node <aa-node2>
```

<aa-node1> and <aa-node2> should be the host names obtained by executing the command “uname –n” on the two assurance nodes.

The recommended configuration is to not allow the failback of resources when the primary node goes up again. This can be changed with using the auto_failback configuration to “on”.

3. /etc/ha.d/haresources

```
<primary node in Assurance cluster> IPaddr::<vip>/24/eth0:0 nfvd
```

The <vip> should be the Virtual IP for the Assurance heartbeat cluster.

2.7 Starting NFVD HA Platform

2.7.1 NFVD GUI VMs

The NFVD GUI components should be started on both the VMs

```
[root@gui-node1 ~]# /opt/HPE/nfvd/bin/nfv-director.sh start
```

The expected output should be:

```
Service Activator (HPSA)              : [Not Installed]
Service Order Smart Adapter (SOSA)    : [Not Installed]
Equipment Connection Pool (ECP)       : [Not Installed]
Lock Manager                          : [Not Installed]
Open Mediation                        : [Not Installed]
SiteScope                             : [Not Installed]
UCA-EBC                               : [Not Installed]
NFVD Assurance Gateway                : [Not Installed]
Apache CouchDB                        : [Started]
UOC Server                            : [Started]
NFVD Authentication Server            : [Started]
NFVD Image Uploader Service           : [Started]
```

```
[root@gui-node2 ~]# /opt/HPE/nfvd/bin/nfv-director.sh start
```

The expected output should be:

```
Service Activator (HPSA)              : [Not Installed]
Service Order Smart Adapter (SOSA)    : [Not Installed]
Equipment Connection Pool (ECP)       : [Not Installed]
Lock Manager                          : [Not Installed]
Open Mediation                        : [Not Installed]
SiteScope                             : [Not Installed]
UCA-EBC                               : [Not Installed]
NFVD Assurance Gateway                : [Not Installed]
Apache CouchDB                        : [Started]
```

```
UOC Server                             : [Started]
NFVD Authentication Server             : [Started]
NFVD Image Uploader Service            : [Started]
```

2.7.2 NFVD Fulfilment VMs

The FF API component is an Active/Active component and should be started on both the Fulfilment VMs

```
[root@ff-node1 ~]# /opt/HPE/nfvd/bin/nfv-director.sh –c activator start
```

The expected output should be:

```
Service Activator (HPSA)               : [Started]
Service Order Smart Adapter (SOSA)     : [Not Running]
Equipment Connection Pool (ECP)        : [Not Running]
Lock Manager                           : [Not Running]
Open Mediation                         : [Not Installed]
SiteScope                              : [Not Installed]
UCA-EBC                                : [Not Installed]
NFVD Assurance Gateway                 : [Not Installed]
Apache CouchDB                         : [Not Installed]
UOC Server                             : [Not Installed]
NFVD Authentication Server             : [Not Installed]
NFVD Image Uploader Service            : [Not Installed]
```

```
[root@ff-node2 ~]# /opt/HPE/nfvd/bin/nfv-director.sh –c activator start
```

The expected output should be:

```
Service Activator (HPSA)               : [Started]
Service Order Smart Adapter (SOSA)     : [Not Running]
Equipment Connection Pool (ECP)        : [Not Running]
Lock Manager                           : [Not Running]
Open Mediation                         : [Not Installed]
SiteScope                              : [Not Installed]
UCA-EBC                                : [Not Installed]
NFVD Assurance Gateway                 : [Not Installed]
Apache CouchDB                         : [Not Installed]
UOC Server                             : [Not Installed]
NFVD Authentication Server             : [Not Installed]
NFVD Image Uploader Service            : [Not Installed]
```

- **The SOSA and LockManager process are Active/Passive and should be started one by one using heartbeat .** Anyway a set of restarts over the same server will be configured to happen (recommended 3) over the same server before switching over another node.

On ff-node1

```
[root@ff-node1 ~]# service heartbeat restart
```

After some time the check the status of processes on this node. Please be patient as hearbeat takes some time to acquire a VIP for the first time.

```
[root@ff-node1 ~]# /opt/HPE/nfvd/bin/nfv-director.sh status
```

The expected output should be:

```
Service Activator (HPSA)              : [Running]
Service Order Smart Adapter (SOSA)    : [Running]
Equipment Connection Pool (ECP)       : [Not Running]
Lock Manager                          : [Running]
Open Mediation                        : [Not Installed]
SiteScope                             : [Not Installed]
UCA-EBC                               : [Not Installed]
NFVD Assurance Gateway                : [Not Installed]
Apache CouchDB                        : [Not Installed]
UOC Server                            : [Not Installed]
NFVD Authentication Server            : [Not Installed]
NFVD Image Uploader Service           : [Not Installed]
```

Note : Please ignore the status of Equipment Connection Pool (ECP) process it's no more used..

On ff-node2

```
[root@ff-node2 ~]# service heartbeat restart
```

After some time the check the status of processes on this node.

```
[root@ff-node2 ~]# /opt/HPE/nfvd/bin/nfv-director.sh status
```

The expected output should be:

```
Service Activator (HPSA)              : [Running]
Service Order Smart Adapter (SOSA)    : [Not Running]
Equipment Connection Pool (ECP)       : [Not Running]
Lock Manager                          : [Not Running]
Open Mediation                        : [Not Installed]
SiteScope                             : [Not Installed]
UCA-EBC                               : [Not Installed]
NFVD Assurance Gateway                : [Not Installed]
Apache CouchDB                        : [Not Installed]
UOC Server                            : [Not Installed]
NFVD Authentication Server            : [Not Installed]
NFVD Image Uploader Service           : [Not Installed]
```

2.7.3 NFVD Sitescope VMs

The SiteScope components should be started on both the VMs if not already running.

```
[root@sis-node1 ~]# /opt/HPE/nfvd/bin/nfv-director.sh status
```

If Sitescope is not running then start the SiteScope process using the below commands.

```
[root@sis-node1 ~]# /opt/HPE/nfvd/bin/nfv-director.sh start
```

The expected output should be:

```
Service Activator (HPSA)            : [Not Installed]
Service Order Smart Adapter (SOSA)  : [Not Installed]
Equipment Connection Pool (ECP)     : [Not Installed]
Lock Manager                        : [Not Installed]
Open Mediation                      : [Not Installed]
SiteScope                           : [Started]
UCA-EBC                             : [Not Installed]
NFVD Assurance Gateway              : [Not Installed]
Apache CouchDB                      : [Not Installed]
UOC Server                          : [Not Installed]
NFVD Authentication Server          : [Not Installed]
NFVD Image Uploader Service         : [Not Installed]
```

```
[root@sis-node2 ~]# /opt/HPE/nfvd/bin/nfv-director.sh -start
```

The expected output should be:

```
Service Activator (HPSA)            : [Not Installed]
Service Order Smart Adapter (SOSA)  : [Not Installed]
Equipment Connection Pool (ECP)     : [Not Installed]
Lock Manager                        : [Not Installed]
Open Mediation                      : [Not Installed]
SiteScope                           : [Started]
UCA-EBC                             : [Not Installed]
NFVD Assurance Gateway              : [Not Installed]
Apache CouchDB                      : [Not Installed]
UOC Server                          : [Not Installed]
NFVD Authentication Server          : [Not Installed]
NFVD Image Uploader Service         : [Not Installed]
```

2.7.4 NFVD Assurance VMs

The Assurance components are Active/Passive and should be started one by one using heartbeat

On aa-node1

```
[root@aa-node1 ~]# service heartbeat restart
```

After some time the check the status of processes on this node. Please be patient as hearbeat takes some time to acquire a VIP for the first time.

```
[root@aa-node1 ~]# /opt/HPE/nfvd/bin/nfv-director.sh status
```

The expected output should be:

```
Service Activator (HPSA)            : [Not Installed]
Service Order Smart Adapter (SOSA)  : [Not Installed]
Equipment Connection Pool (ECP)     : [Not Installed]
Lock Manager                        : [Not Installed]
Open Mediation                      : [Running]
SiteScope                           : [Not Installed]
UCA-EBC                             : [Running]
NFVD Assurance Gateway              : [Running]
Apache CouchDB                      : [Not Installed]
UOC Server                          : [Not Installed]
NFVD Authentication Server          : [Not Installed]
NFVD Image Uploader Service         : [Not Installed]
```

On aa-node2

```
[root@aa-node2 ~]# service heartbeat restart
```

After some time the check the status of processes on this node.

```
[root@aa-node2 ~]# /opt/HPE/nfvd/bin/nfv-director.sh status
```

The expected output should be:

```
Service Activator (HPSA)            : [Not Installed]
Service Order Smart Adapter (SOSA)  : [Not Installed]
Equipment Connection Pool (ECP)     : [Not Installed]
Lock Manager                        : [Not Installed]
Open Mediation                      : [Not Running]
SiteScope                           : [Not Installed]
UCA-EBC                             : [Not Running]
NFVD Assurance Gateway              : [Not Running]
Apache CouchDB                      : [Not Installed]
UOC Server                          : [Not Installed]
NFVD Authentication Server          : [Not Installed]
NFVD Image Uploader Service         : [Not Installed]
```

2.8 Load Balancer Configurations

This section describes the health check monitoring tools available on the NFVD GUI, Fulfillment API, and SiteScope components. It then describes how these monitoring tools can be used by an external load balancer to provide Active/Active load balancing and High Availability for the NFVD GUI and Fulfillment API components, and Active/Hot Standby High Availability for SiteScope.

Example configurations for each component are given for HAProxy and F5 load balancers.

2.8.1 NFVD GUI

The NFVD GUI components operate in an Active/Active mode, with requests sent to each GUI instance in a round-robin fashion by an external load balancer. The GUI components also require access to shared storage in order to support the storage of VNF images

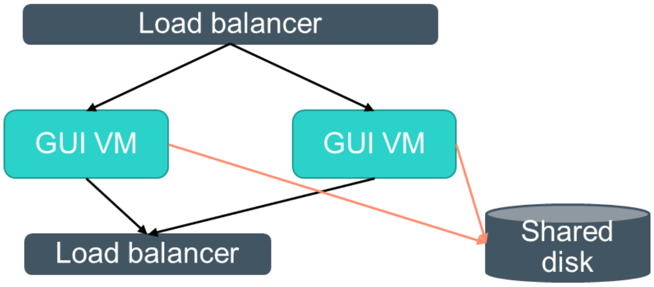

Figure 8: NFVD GUI Load Balancer / Shared Storage configuration

2.8.1.1 Health Check Script

Each NFVD GUI instance has a health check script:

```
/opt/uoc2/scripts/monitor
```

This script is registered with the xinitd daemon by the HA installation script. It can be invoked by an external load balancer with an “http GET” on port 3001.

When invoked using http GET the output would be.

```
# curl -X GET  http://<GUI-node1-IP>:3001
```

Expected result when NFVD GUI instance is up.

```
NFVD-GUI is running
```

Expected result when NFVD GUI instance is down.

```
NFVD-GUI Unavailable
```

2.8.1.2 Load Balancer Configuration

An external load balancer is required to provide High Availability for the NFVD GUI component. The minimal requirements for the load balancer are as follows:

The load balancer is has network connectivity to the NFVD GUI servers and clients

- The load balancer must be able to provide “sticky” sessions, so that end users are consistently connected to the same GUI server.
- The load balancer must provide the ability to re-dispatch existing connections on a given server to another server in case of a failure
- The load balancer must be able to use a periodic http or https health monitor to check the two GUI servers, and perform a failover if the health check fails.
- The load balancer must be able to do a failover if the connection to a GUI server is lost for more than a specific timeout period

A typical http/https health monitor could be configured as follows:

- periodically send the following http/https GET strings:
    - GET http://GUI_node1:3001
    - GET http://GUI_node2:3001
- The health check period can be set to 15 seconds
- The connectivity timeout period can be set to 46 seconds

2.8.1.3 HA_Proxy NFVD GUI Load Balancer configuration

An example of the NFVD GUI portion of a HAProxy configuration file is as follows (a complete haproxy.cfg file is presented in Section 2.8.4

```
# GUI LB VIP
listen  http web gui 16.16.88.154:3000
        mode http
        balance roundrobin  # Load Balancing algorithm : Active/Active
        option httpchk GET
        option forwardfor
        option redispatch       #Send connections to the other server in case of failure
        server server1 16.16.88.201:3000 weight 1 maxconn 512 cookie GUI1 check port 3001 inter
15s
        server server2 16.16.88.237:3000 weight 1 maxconn 512 cookie GUI2 check port 3001 inter
15s
#       Health check is: GET http:<GUI_IP>:3001, sent on port 3001 at an interval of 15 seconds


# GUI SAML LB VIP
listen  http_web_saml_gui 16.16.88.154:38080
        mode http
        balance roundrobin  # Load Balancing algorithm
        option httpchk GET
        option forwardfor
        option redispatch
        server server1 16.16.88.201:38080 weight 1 maxconn 512 cookie GUI1 check port 3001
inter 15s
        server server2 16.16.88.237:38080 weight 1 maxconn 512 cookie GUI2 check port 3001
inter 15s

# GUI IMAGE REPOSITORY ENDPOINT
listen  http_image_repository 16.16.88.154:1337
        mode http
        balance roundrobin  # Load Balancing algorithm
        option httpchk GET
        option forwardfor
        option redispatch
        server server1 16.16.88.201:1337 weight 1 maxconn 512 cookie GUI1 check port 3001 inter
15s
```

```
        server server2 16.16.88.237:1337 weight 1 maxconn 512 cookie GUI2 check port 3001 inter
15s
```

2.8.1.4 F5 NFVD GUI Load Balancer configuration

The following steps should be followed to configure an F5 LTM load balancer for the NFVD GUI:

- Define a pool for the GUI consisting of the two NFVD GUI servers. Assign the GUI VIP to this pool.
- Define an http health monitor for the GUI pool with the following attributes:
    - Interval: 5 seconds
      (time between health check monitor calls to the pool members)
    - Timeout 16 seconds
      (timeout if pool member does not respond with this time. F5 recommendation for this value is 3*Interval + 1)
    - Send string: GET /\r\n
      (http GET string sent to a pool member)
    - Receive string: HTTP 1.1 200 NFVD-UI OK
      (expected health response)
    - Receive Disable string: HTTP 1.1 503 NFVD-UI
      (expected health response on failure)
    - Alias Service Port: 3001
      (port used by the Health Monitor)
- Set the Load Balance method to RoundRobin

For more details, please refer to Section 2.8.5, "Configuring an F5 LTM Load Balancer"

2.8.2 NFVD Fulfillment API

The NFVD Fulfillment API components operate in an Active/Active mode, with requests sent to each Fulfillment API instance in a round-robin fashion by an external load balancer. The Fulfillment API components also require access to shared storage in order to support the storage of VNF images

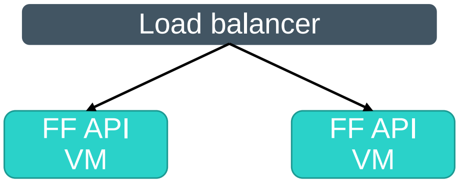

Figure 9: NFVD Fulfillment API Load Balancer configuration

2.8.2.1 Health Check Script

Each instance provides a Health check tool to check the status of NFVD Fulfilment API component. It can be invoked by an external load balancer with an "http GET" on port 8080.

When invoked using http GET the output would be.

```
# curl -X GET  http://<FF-node1-IP>:8080/nfvd-ext/status
```

Expected result when NFVD FF API instance is up.

```
[NFVD FF API OK, NFVD FF LDAP OK, NFVD FF DB OK]
```

Expected result when NFVD FF API instance is down.

```
HTTP/1.1 500 Internal Server Error
```

2.8.2.2 Load Balancer Configuration

An external load balancer is required to provide High Availability for the NFVD Fulfillment API component. The minimal requirements for the load balancer are as follows:

The load balancer is has network connectivity to the NFVD Fulfillment API servers and clients

- The load balancer must be able to provide “sticky” sessions, so that end users are consistently connected to the same Fulfillment API server.
- The load balancer must provide the ability to re-dispatch existing connections on a given server to another server in case of a failure
- The load balancer must be able to use a periodic http or https health monitor to check the two Fulfillment API servers, and perform a failover if the health check fails.
- The load balancer must be able to do a failover if the connection to a Fulfillment API server is lost for more than a specific timeout period

A typical http/https health monitor could be configured as follows:

- periodically send the following http/https GET strings:
    - GET http://FF_API _node1:8080/nfvd-ext/status
    - GET http://FF_API _node2: 8080/nfvd-ext/status

- The health check period can be set to 15 seconds
- The connectivity timeout period can be set to 46 seconds

2.8.2.3 HA_Proxy Fulfillment API Load Balancer configuration

An example of the Fulfillment API portion of a HAProxy configuration file is as follows (a complete haproxy.cfg file is presented in Section 2.8.4

```
# FF API LB VIP
listen  http_web_ff 16.16.88.156:8080
        mode http
        balance roundrobin  # Load Balancing algorithm : Active/Active
        option httpchk GET /nfvd-ext/status
        option forwardfor
        option redispatch       #Send connections to the other server in case of failure
        server server1 16.16.89.225 weight 1  maxconn 512 cookie FF1 check port 8080 inter 15s
        server server2 16.16.88.135 weight 1  maxconn 512 cookie FF2 check port 8080 inter 15s
#       Health check is GET  http:// <FFHostIP>:8080/nfvd-ext/status sent on port 8080 at an
interval of 15 seconds
```

2.8.2.4 F5 Fulfillment API Load Balancer configuration

The following steps should be followed to configure an F5 LTM load balancer for the NFVD GUI:

- Define a pool for the GUI consisting of the two NFVD Fulfillment API servers. Assign the Fulfillment API VIP to this pool.
- Define an http health monitor for the GUI pool with the following attributes:
  - Interval: 5 seconds
    (time between health check monitor calls to the pool members)
  - Timeout 16 seconds
    (timout if pool member does not respond with this time. F5 recommendation for this value is 3*Interval + 1)
  - Send string: GET /nfvd-ext/status\r\n
    (http GET string sent to a pool member)
  - Alias Service Port: 8080
    (port used by the Health Monitor)
- Set the Load Balance method to RoundRobin

For more details, please refer to Section 2.8.5, "Configuring an F5 LTM Load Balancer"

2.8.3 SiteScope

HPE SiteScope operates in a Primary/Hot Standby mode. An external load balancer is used to direct traffic to the Primary node. If ever a failure is detected on the Primary node, then the Hot Standby secondary node takes over monitoring, and all traffic is directed to the Hot Standby. When the Primary node recovers, the traffic is then re-routed back to the Primary node.

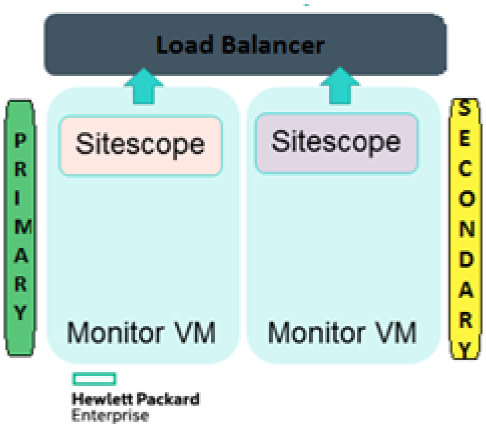

Figure 10: SiteScope Load Balancer Configuration

2.8.3.1 Health Check Script

Each SiteScope instance has a health check script:

```
/opt/HPE/nfvd/solutions/ha-example/sis_check_primary.sh
```

This script has to be registered with the xinitd daemon so that it can be invoked by an external load balancer with an "http GET" on port 8898.

On each of the SiteScope VMs configure the Sitescope health Service as described below.

Edit /etc/services to have the below config at the end of the file, save and quit

```
sischkprm       8898/tcp                 # sischkprm
```

Edit/ Create file - /etc/xinetd.d/sischkprm to have the below config at the end of the file, save and quit.

```
# default: on
# description: sischkprm
service sischkprm
{
        flags           = REUSE
        socket_type     = stream
        port            = 8898
        wait            = no
        user            = root
        server          = /opt/HPE/nfvd/solutions/ha-example/sis_check_primary.sh
        log_on_failure  += USERID
        disable         = no
        only_from       = 0.0.0.0/0
        per_source      = UNLIMITED
}
```

Restart xinetd

```
/etc/init.d/xinetd restart
```

The Sitescope health script is ready now and when invoked using http GET the output would be.

```
# curl -X GET  http://<SIS-node1-IP>:8898
```

Expected result when SiteScope instance is up.

```
SiteScope is running.
```

Expected result when SiteScope instance is down.

```
SiteScope is down.
```

2.8.3.2 Load Balancer Configuration

An external load balancer is required to provide High Availability for SiteScope. The minimal requirements for the load balancer are as follows:

The load balancer is has network connectivity to the SiteScope servers and clients

- The load balancer must be able to provide “sticky” sessions, so that end users are consistently connected to the same SiteScope server.
- The load balancer must provide the ability to re-dispatch existing connections on a given server to another server in case of a failure
- The load balancer must be able to use a periodic http or https health monitor to check the two SiteScope servers, and perform a failover from the primary to the secondary if the health check fails on the primary.
- When the hot standby secondary server is active, and the primary server comes back on line, the load balancer must re-direct traffic back to the primary
- The load balancer must be able to do a failover if the connection to a SiteScope server is lost for more than a specific timeout period

A typical http/https health monitor could be configured as follows:

- periodically send the following http/https GET strings:
  - GET http://<SiteScopePrimaryIP>:8898
  - GET http://<SitescopeFailoverIP>:8898
- The health check period can be set to 15 seconds
- The connectivity timeout period can be set to 46 seconds

2.8.3.3 HA_Proxy SiteScope Load Balancer configuration

An example of the NFVD GUI portion of a HAProxy configuration file is as follows (a complete haproxy.cfg file is presented in Section 2.8.4

```
listen  http_web_sis 16.16.95.21:18888
        mode http
        balance roundrobin  # Load Balancing algorithm : Active / hot standby
        option httpchk GET
        option forwardfor
        option redispatch       #Send connections to the other server in case of failure
        server server1 16.16.88.203:18888 weight 1 maxconn 512 cookie SiS1 check port 8898
inter 15s fall 1 rise 3
        server server2 16.16.88.182:18888 weight 1 maxconn 512 cookie SiS2 check port 8898
inter 15s backup
#       Health check is: GET http:<SiteScope IP>:8898, sent on port 8898 at an interval of 15
seconds
#       First server is primary - it is taken offline after 1 failure, and takes over again
after 3 successes (45 seconds)
#       Second server is backup, and takes over only when first server fails
```

2.8.3.4 F5 NFVD SiteScope Load Balancer configuration

The following steps should be followed to configure an F5 LTM load balancer for the NFVD GUI:

- Define a pool for the GUI consisting of the two NFVD SiteScope servers. Assign the SiteScope VIP to this pool.
- Assign a priority group to each pool member. Suggested values are 10 for the primary, and 5 for the failover server.
-
- Define an http health monitor for the SiteScope pool with the following attributes:
  - Interval: 5 seconds
    (time between health check monitor calls to the pool members)
  - Timeout 16 seconds
    (timout if pool member does not respond with this time. F5 recommendation for this value is 3*Interval + 1)
  - Send string: GET /\r\n
    (http GET string sent to a pool member)
  - Receive string:HTTP 1.1 200 NFVD SiteScope OK
    (expected health response)
  - Receive Disable string: HTTP 1.1 503 NFVD
    (expected health response on failure)
  - Alias Service Port: 3001
    (port used by the Health Monitor)
- Set the Load Balance method to RoundRobin
- Set the Priority Group Activation to “Less than” “1”
  This will guarantee that the primary server receives traffic when it is up, and the lower-priority failover server only receives traffic when the primary is down.

For more details, please refer to Section 2.8.5, "Configuring an F5 LTM Load Balancer"

2.8.4 An Example HAProxy Load Balancer Configuration Script

```
#---------------------------------------------------------------------
# Example configuration for the NFVD HAProxy load balancer
#---------------------------------------------------------------------

#---------------------------------------------------------------------
# Global settings
#---------------------------------------------------------------------
global
    # to have these messages end up in /var/log/haproxy.log you will
    # need to:
    #
    # 1) configure syslog to accept network log events.  This is done
    #    by adding the '-r' option to the SYSLOGD_OPTIONS in
    #    /etc/sysconfig/syslog
    #
    # 2) configure local2 events to go to the /var/log/haproxy.log
    #   file. A line like the following can be added to
    #   /etc/sysconfig/syslog
    #
    #    local2.*                       /var/log/haproxy.log
    #
    log         127.0.0.1 local2

    chroot      /var/lib/haproxy
    pidfile     /var/run/haproxy.pid
    maxconn     4000
    user        haproxy
    group       haproxy
    daemon

    # turn on stats unix socket
    stats socket /var/lib/haproxy/stats

#---------------------------------------------------------------------
# common defaults that all the 'listen' and 'backend' sections will
# use if not designated in their block
#---------------------------------------------------------------------
defaults
    mode                    http
    log                     global
    option                  httplog
    option                  dontlognull
    option http-server-close
    option forwardfor       except 127.0.0.0/8
    option                  redispatch
    retries                 3
    timeout http-request    10s
    timeout queue           1m
    timeout connect         10s
    timeout client          30s
    timeout server          30s
    timeout http-keep-alive 10s
    timeout check           10s
    maxconn                 3000

       #      Use cookies to ensure "sticky" sessions
       cookie NFVD_SRV insert
```

```
# GUI LB VIP
listen  http_web_gui 16.16.88.154:3000
        mode http
        balance roundrobin  # Load Balancing algorithm : Active/Active
        option httpchk GET
        option forwardfor
       option redispatch   #Send connections to the other server in case of failure
        server server1 16.16.88.201:3000 weight 1 maxconn 512 cookie GUI1 check port 3001 inter 
15s
        server server2 16.16.88.237:3000 weight 1 maxconn 512 cookie GUI2 check port 3001 inter 
15s
#      Health check is: GET http:<GUI_IP>:3001, sent on port 3001 at an interval of 15 seconds


# GUI SAML LB VIP
listen  http web saml gui 16.16.88.154:38080
        mode http
        balance roundrobin  # Load Balancing algorithm
        option httpchk GET
        option forwardfor
        option redispatch
        server server1 16.16.88.201:38080 weight 1 maxconn 512 cookie GUI1 check port 3001 
inter 15s
        server server2 16.16.88.237:38080 weight 1 maxconn 512 cookie GUI2 check port 3001 
inter 15s

# GUI IMAGE REPOSITORY ENDPOINT
listen  http image repository 16.16.88.154:1337
        mode http
        balance roundrobin  # Load Balancing algorithm
        option httpchk GET
        option forwardfor
        option redispatch
        server server1 16.16.88.201:1337 weight 1 maxconn 512 cookie GUI1 check port 3001 inter 
15s
        server server2 16.16.88.237:1337 weight 1 maxconn 512 cookie GUI2 check port 3001 inter 
15s

# FF API LB VIP
listen  http web ff 16.16.88.156:8080
        mode http
        balance roundrobin  # Load Balancing algorithm : Active/Active
        option httpchk GET /nfvd-ext/status
        option forwardfor
        option redispatch       #Send connections to the other server in case of failure
        server server1 16.16.89.225 weight 1  maxconn 512 cookie FF1 check port 8080 inter 15s
        server server2 16.16.88.135 weight 1  maxconn 512 cookie FF2 check port 8080 inter 15s
#       Health check is GET  http:// <FFHostIP>:8080/nfvd-ext/status sent on port 8080 at an 
interval of 15 seconds

# SiteScope  LB VIP
listen  http_web_sis 16.16.95.21:18888
        mode http
        balance roundrobin  # Load Balancing algorithm : Active / hot standby
        option httpchk GET
        option forwardfor
        option redispatch       #Send connections to the other server in case of failure
        server server1 16.16.88.203:18888 weight 1 maxconn 512 cookie SiS1 check port 8898 
inter 15s fall 1 rise 3
        server server2 16.16.88.182:18888 weight 1 maxconn 512 cookie SiS2 check port 8898 
inter 15s backup
#       Health check is: GET http:<SiteScope_IP>:8898, sent on port 8898 at an interval of 15 
seconds
```

```
#        First server is primary - it is taken offline after 1 failure, and takes over again
after 3 successes (45 seconds)
#        Second server is backup, and takes over only when first server fails

#STATS
listen stats *:1936
stats enable
stats uri /
stats hide-version
stats auth udg:udg
```

2.8.5 Configuring an F5 LTM Load Balancer

First, define pools for each of the NFVD GUI, Fulfillment API< and SiteScope components. Here is an example of configuring the GUI pool:

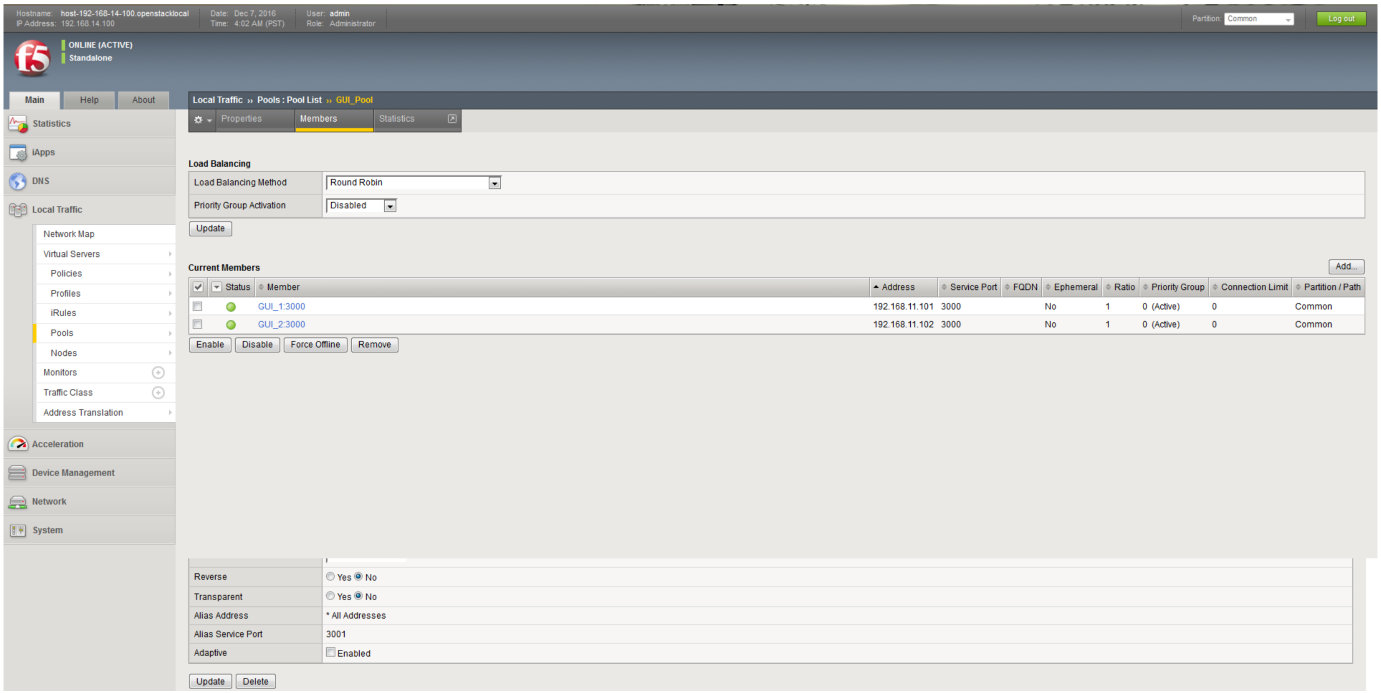

Local Traffic » Pools : Pool List » GUI_Pool

Load Balancing

Load Balancing Method	Round Robin
Priority Group Activation	Disabled

Current Members

Status	Member	Address	Service Port	FQDN	Ephemeral	Ratio	Priority Group	Connection Limit	Partition / Path
●	GUI_1:3000	192.168.11.101	3000		No	1	0 (Active)	0	Common
●	GUI_2:3000	192.168.11.102	3000		No	1	0 (Active)	0	Common

Reverse	○ Yes ● No
Transparent	○ Yes ● No
Alias Address	* All Addresses
Alias Service Port	3001
Adaptive	□ Enabled

Figure 11: F5 Configuration for the GUI pool

Next, define the http health monitors for each pool:

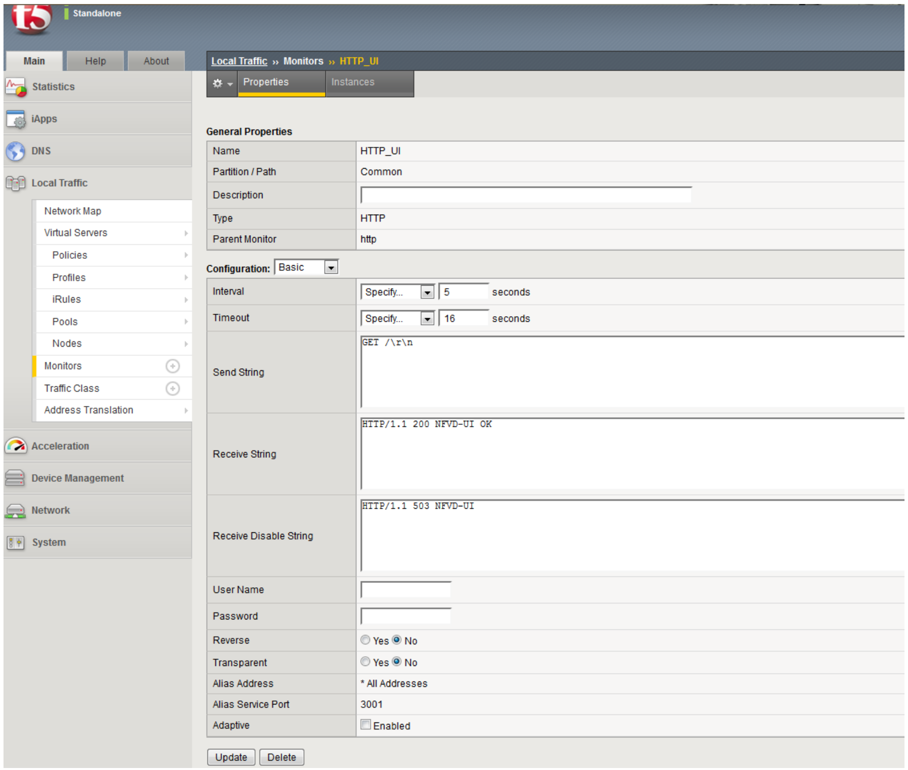

Local Traffic » Monitors » HTTP_UI

General Properties

Name	HTTP_UI
Partition / Path	Common
Description	
Type	HTTP
Parent Monitor	http

Configuration: Basic

Interval	Specify... 5 seconds
Timeout	Specify... 16 seconds
Send String	GET /\r\n
Receive String	HTTP/1.1 200 NFVD-UI OK
Receive Disable String	HTTP/1.1 503 NFVD-UI
User Name	
Password	
Reverse	○ Yes ● No
Transparent	○ Yes ● No
Alias Address	* All Addresses
Alias Service Port	3001
Adaptive	□ Enabled

Figure 12: F5 http health monitor for the GUI pool

Assign the health monitor to the appropriate pool:

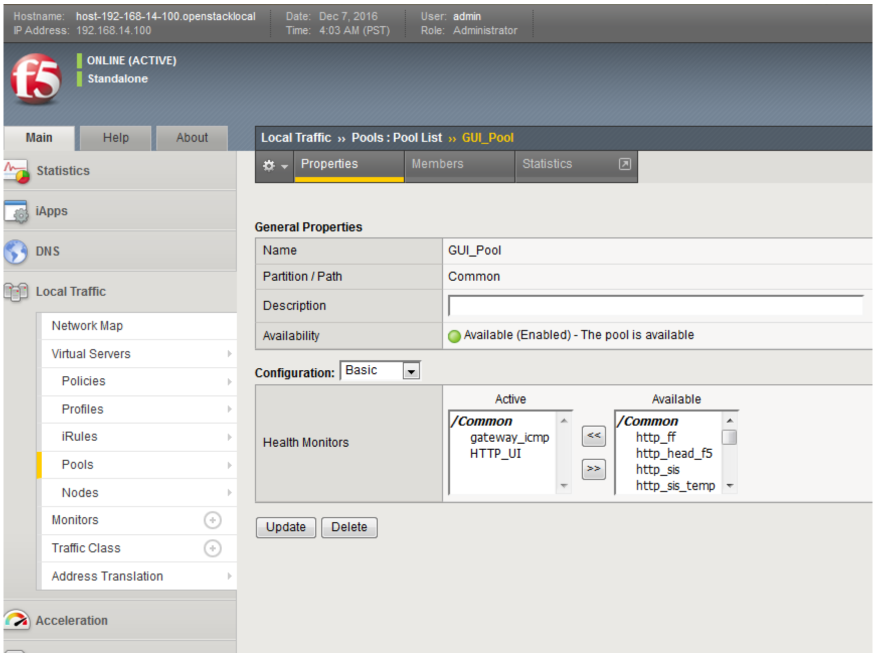

Hostname: host-192-168-14-100.openstacklocal
IP Address: 192.168.14.100
Date: Dec 7, 2016
Time: 4:03 AM (PST)
User: admin
Role: Administrator

ONLINE (ACTIVE)
Standalone

Local Traffic » Pools : Pool List » GUI_Pool

General Properties

Name	GUI_Pool
Partition / Path	Common
Description	
Availability	Available (Enabled) - The pool is available

Configuration: Basic

	Active	Available
Health Monitors	/Common gateway_icmp HTTP_UI	/Common http_ff http_head_f5 http_sis http_sis_temp

Figure 13: Assign the http GUI health monitor to the GUI pool

Follow the same procedure for the Fulfillment API pool and the SiteScope pool. The only difference is that SiteScope needs to be set up as Active / Hot Standby, so the failover server only receives traffic when the primary is down. This is accomplished by assigning a priority to each server (10 to the primary, 5 to the failover). The active server with highest priority will receive traffic. Then set the Priority Group Activation to “Less than…” “1”:

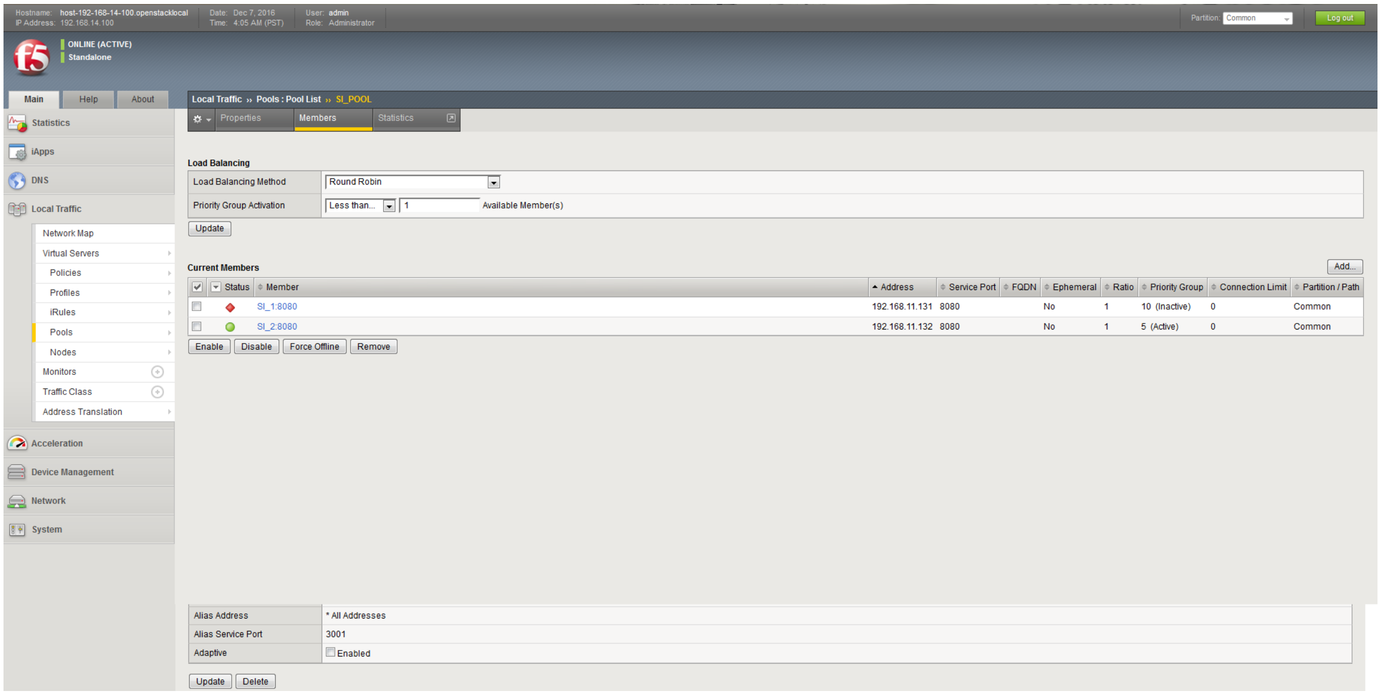

Hostname: host-192-168-14-100.openstacklocal
IP Address: 192.168.14.100
Date: Dec 7, 2016
Time: 4:05 AM (PST)
User: admin
Role: Administrator

Local Traffic » Pools : Pool List » SI_POOL

Load Balancing

Load Balancing Method	Round Robin
Priority Group Activation	Less than... 1 Available Member(s)

Update

Current Members

Add...

☑	Status	Member	Address	Service Port	FQDN	Ephemeral	Ratio	Priority Group	Connection Limit	Partition / Path
□	◆	SI_1:8080	192.168.11.131	8080		No	1	10 (Inactive)	0	Common
□	●	SI_2:8080	192.168.11.132	8080		No	1	5 (Active)	0	Common

Enable Disable Force Offline Remove

Alias Address	* All Addresses
Alias Service Port	3001
Adaptive	□ Enabled

Update Delete

Figure 14: Priority Group Setup for the SiteScope Pool

2.9 Configuring Monitoring Tools

In addition to the health monitoring tools that can be called from an external load balancer, there are local monitoring tools that can be called from *cron* to check the health of processes NFVD processes and then take appropriate remedial actions automatically.

2.9.1 Configuring Fulfilment VM Monitoring Tools

The HA monitoring scripts are located under the folder /opt/HPE/nfvd/fulfillment/scripts.

All the configuration parameters changeable for a script are available as a variable in the header of the script. There are two pairs of scripts, one pair for Jboss and the other for SOSA/Lock Manager

Jboss scripts:

- nfvd_jboss_restart.sh:
  - this script can stop and start the Jboss instance of a machine
  - By default, it logs to /var/opt/OV/ServiceActivator/log/nfvd_checker/nfvd_jboss_restart.log

- nfvd_jboss_check.sh:
  - this script is responsible to check the existence of the Jboss process that is serving both MWFM and FF API.
  - If the jboss process is not present, it invokes the restart script to try to start the Jboss process, maximum 3 times.
  - After the third time, it logs a failure and it will not try to start Jboss anymore.

  - By default, it logs to/var/opt/OV/ServiceActivator/log/nfvd_checker/nfvd_jboss_check.log. Number of retries will be stored in /var/opt/OV/ServiceActivator/tmp/jboss_retries by default.

SOSA / Lock Manager scripts:

- nfvd_sosa_lm_restart.sh:
  - this script can stop and start SOSA and LockManager locally in the machine.
  - By default, it logs to /var/opt/OV/ServiceActivator/log/nfvd_checker/nfvd_sosa_lm_restart.log.
  - As by default a set of restarts over the same server will be configured to happen (recommended 3) over the same server before switching over another node.

- nfvd_sosa_lm_check.sh:
  - this script is responsible to check the existence of the SOSA and LockManager processes in an active/passive HA configuration.
  - If the any of the processes are not present, it will invoke the restart script to try to start them, maximum 3 times.
  - After the third time, it logs a failure and it will not try to start Jboss anymore.
  - By default, it logs to /var/opt/OV/ServiceActivator/log/nfvd_checker/ nfvd_sosa_lm_check.log.
  - Number of retries will be stored in /opt/OV/ServiceActivator/EP/SOSA/tmp/sosa_retries for SOSA, and /opt/OV/ServiceActivator/EP/LockManager/tmp/lock_manager_retries for LockManager, by default.

2.9.1.1 Configuration

Note : The configuration below is to be done on both Fulfillment VMs

Edit crontab and add the following lines:

```
*/5  *   *   *   * /opt/HPE/nfvd/fulfillment/scripts/nfvd_sosa_lm_check.sh > /dev/null 2>&1
*/5  *   *   *   * /opt/HPE/nfvd/fulfillment/scripts/nfvd_jboss_check.sh > /dev/null 2>&1
```

Next, set execution permissions on these scripts:

```
chmod +x /opt/HPE/nfvd/fulfillment/scripts/*.sh
```

2.9.2 Configuring SiteScope VM Monitoring tool

sis_check.sh (located in /opt/HPE/nfvd/solutions/ha-example) is an independent script registered to cron and triggered at regular intervals to check and ensure that the Sitescope process is OK on Primary and Secondary Node.

The tool does not display anything on the standard output. Output is logged in /var/opt/HPE/nfvd/halogs/nfvdha-sis.log.

It returns 0 in case of success and 1 in case of error.

2.9.2.1 Configuration.

Note : The configuration below is to be done on both SiteScope VMs

Edit crontab and add the following line:

```
*/5  *   *   *   * /opt/HPE/nfvd/solutions/ha-example/sis_check.sh localhost > /dev/null 2>&1
```

2.9.3 Configuring Assurance VM Monitoring tool.

The script aa_check.sh (located in /opt/HPE/nfvd/solutions/ha-example) is an independent script registered to cron and triggered at regular intervals to check the process status. It has to be used with the Load balancer.

2.9.3.1 Configuration

Note : The below configuration is to be done on both Assurance VMs

Edit crontab and add the following line:

```
*/5  *  *  *  * /opt/HPE/nfvd/solutions/ha-example/aa_check.sh   <AGW_VIP> > /dev/null 2>&1
```

Note: the AGW_VIP should provide the actual IP Address (i.e x.x.x.x) instead of hostname.

2.10 Validation of the NFVD HA setup

This chapter provides procedures to validate whether the HA installation of NFVD has been done successfully.

2.10.1 Validation of the NFVD GUI VMs

Open a browser to the address http://<GUI-VIP>:3000   The NFVD GUI login page should appear, and it should be possible to log into NFVD

Repeat the above step multiple times to ensure that every time it's successful.

This confirms that:

- The external Load balancer is configured correctly to redirect to NFVD GUI VMs
- The NFVD GUI VMs are running properly.

2.10.2 Validation of Image Management

Prerequisite: ensure that there is a Cirros image available

First log on to NFVD as a Domain user:

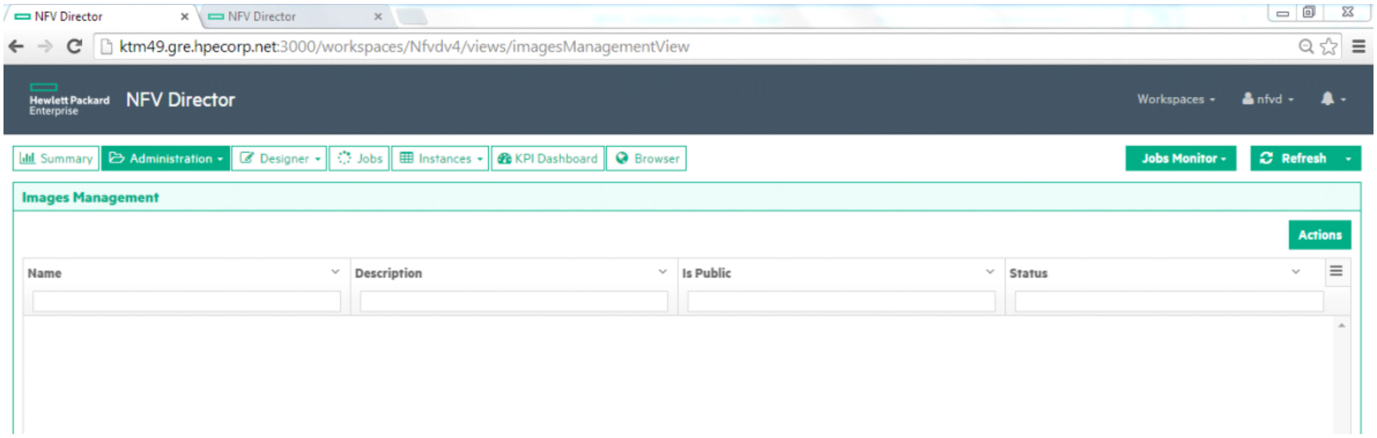

Figure 15: NFVD Domain user logon

Then:

- Select Administration --> Images Management menu.
- Click on Actions --> Create Image
- Fill in the Image screen as follows

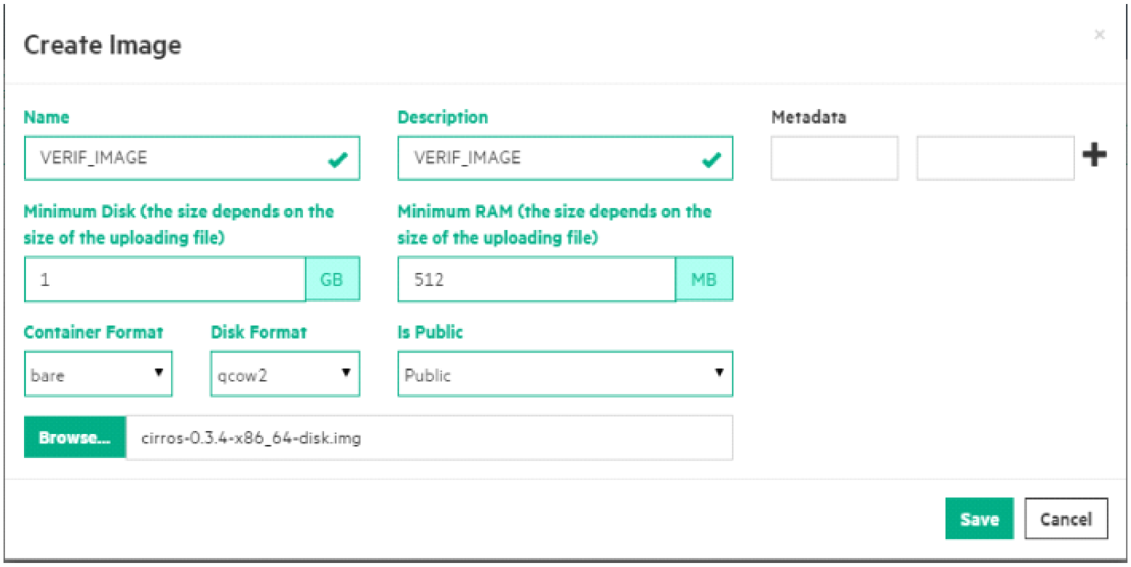

Create Image

Name: VERIF_IMAGE	Description: VERIF_IMAGE	Metadata
Minimum Disk (the size depends on the size of the uploading file): 1 GB	Minimum RAM (the size depends on the size of the uploading file): 512 MB	
Container Format: bare / Disk Format: qcow2	Is Public: Public	

Browse... cirros-0.3.4-x86_64-disk.img

Figure 16: NFVD Image creation

- Browse to the cirros image on your PC
- Click on “Save”

At this point, there should be a progress bar that starts to move under the “Status”

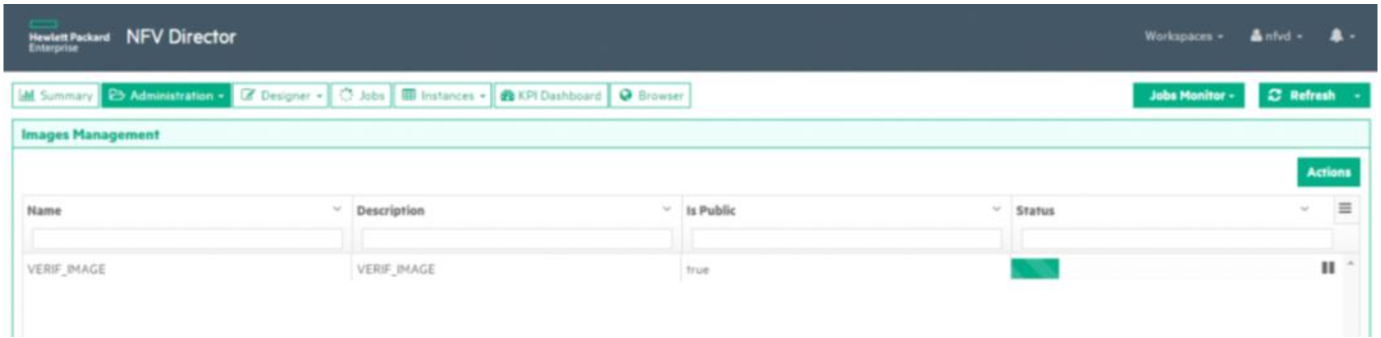

Name	Description	Is Public	Status
VERIF_IMAGE	VERIF_IMAGE	true	

Figure 17: NFVD image upload underway

Wait for the upload to be finished:

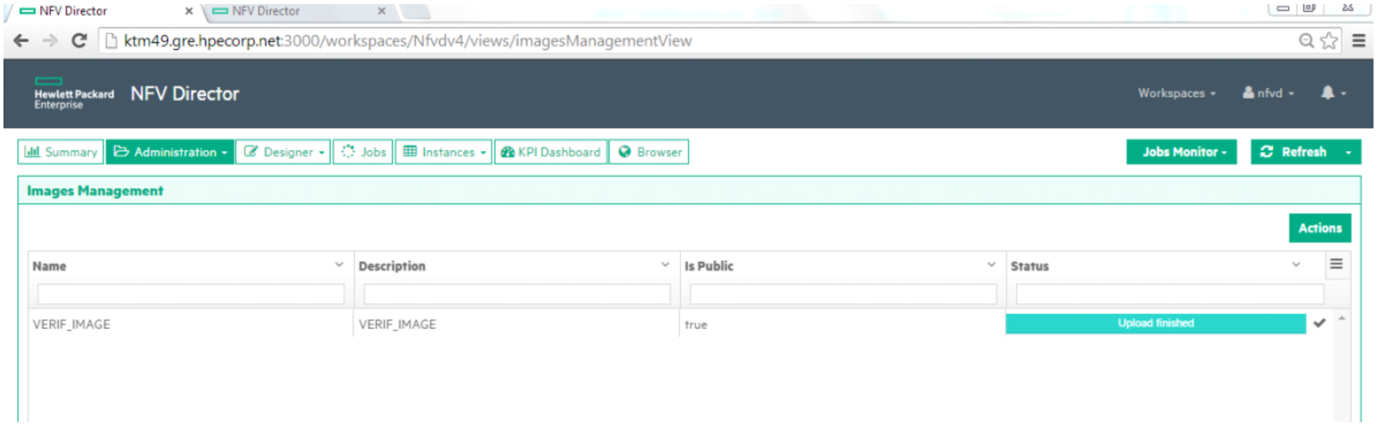

Name	Description	Is Public	Status
VERIF_IMAGE	VERIF_IMAGE	true	Upload finished ✓

Figure 18: NFVD Image Upload Complete

2.10.3 Validation of the Fulfillment VMs

2.10.3.1 FF API component.

Open a browser to the address http://<FF-API-VIP>:8080/activator The activator login page should appear, and it should be possible to log into activator.

Repeat the above step multiple times to ensure that every time it's successful.

This confirms that:

- The external Loadbalancer is configured correctly to redirect to NFVD FF API components
- The NFVD API components are running properly.

2.10.3.2 SOSA and LockManager components.

On the Fulfilment VMs, execute the following command:

```
[root@ff-node1 ~]# cl_status rscstatus
```

The expected result is:

```
all
```

This confirms that ff-node1 is Active and is in control of resources

```
[root@ff-node1 ~]# /etc/init.d/activator-ep check
```

The expected result is:

```
HP Service Activator SOSA application is running
HP Service Activator Lock Manager application is running
```

Do a Manual Switchover of ff-node1 to ff-node2

```
[root@ff-node1 ~]# /usr/share/heartbeat/hb_standby all
```

The expected result is:

```
Going standby [all].
```

After some time, check on ff-node2

```
[root@ff-node2 ~]# cl_status rscstatus
```

The expected result is:

```
all
```

```
[root@ff-node2 ~]# /etc/init.d/activator-ep check
```

The expected result is:

```
HP Service Activator SOSA application is running
HP Service Activator Lock Manager application is running
```

The ff-node2 has now become Active node and ff-node1 is now Passive w.r.t to SOSA and LockManager components.

2.10.4 Validation of the Assurance VMs

On the Assurance VMs, execute the following command:

```
[root@aa-node1 ~]# cl_status rscstatus
```

The expected result is:

```
all
```

This confirms that aa-node1 is Active and is in control of resources

```
[root@aa-node1 ~]# /etc/init.d/nfvd status
```

The expected result is:

```
Status of NFVD processes
NFVD processes are running
```

Do a Manual Switchover of aa-node1 to aa-node2

```
[root@aa-node1 ~]# /usr/share/heartbeat/hb_standby all
```

The expected result is:

```
Going standby [all].
```

After some time, check on aa-node2

```
[root@aa-node2 ~]# cl_status rscstatus
```

The expected result is:

```
all
```

```
[root@aa-node2 ~]# /etc/init.d/nfvd status
```

The expected result is:

```
Status of NFVD processes
NFVD processes are running
```

2.10.5 Validation of the SiteScopeVMs

Open a browser to the address http://<sis-vip>:18888/sitescope, login page should appear, and it should be possible to log into SiteScope

Repeat the above step multiple times to ensure that every time it is successful.

This check confirms that:

- The external load balancer is configured correctly to direct traffic to the NFVD SiteScope VMs
- The SiteScope VMs are running properly.

2.10.6 Validation of the NFVD Self Monitors.

Open a browser to the address http://<sis-vip>:18888/SiteScope/MultiView then login and you should see NFVD Self Monitors Running as shown in the below screen shot.

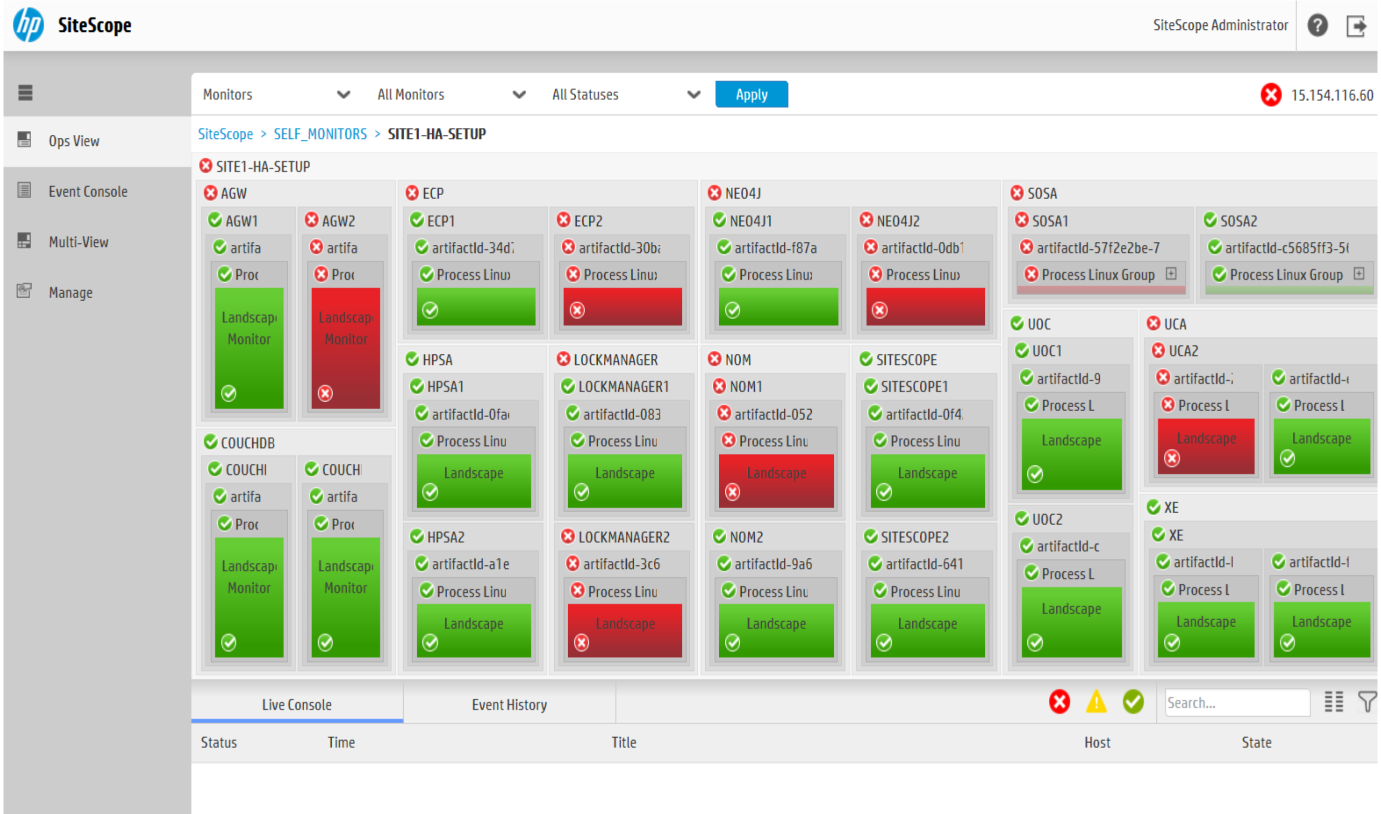

Note : Some monitors in the above GUI are in Red colour because these components are in Passive state as per se the NFVD HA architecture.

Chapter 3
Integrating with the VIM to discover resources.

Once the HA setup is validated follow “NFV Director V4.2 Edition 1 VIM Integration Guide” to trigger the discovery process.

Note in the “NFV Director V4.2 Edition 1 VIM Integration Guide” wherever there is a mention of <FF_HOST> hostname or ip address, replace it with <FF-VIP> the one that is represented by external load balancer.

For example.

In “**Chapter 3** Triggering Discovery process”

Example:

```
./nfvd_createVIM.sh -host 16.16.88.156 -port 8080 -vimname FC33 -url https://10.207.114.100:5000/v2.0/tokens -username
admin -password YMCtbGCT9 -tenantname admin -vimcategory HELION -vimtype HCG
```

Here 16.16.88.156 is he FF-VIP represented by an external load balancer.

The command nfvd_createVIM.sh should be run on the AA node where OpenMediation is up and running.

You can use the command nfv-director.sh –c openmediation on each AA VM to know where OpenMediation is running.

Once the discovery process is done, login into the NFVD GUI http://<GUI-VIP>:3000 to see the discovered resources. **Error! Hyperlink reference not valid.**

To disable discovery you can run the following command from any of the AA VM.

```
sh disable_discovery.sh –h <AA-VIP> -m https
```

To enable discovery you can run the following command from any of the AA VM.

```
sh enable_discovery.sh –h <AA-VIP> -m https
```

The system is now ready for VNF deployments.

Chapter 4 Administration of NFVD HA VMs

4.1 Start/Stop/Status of NFVD VMs

To start the NFVD VMs in HA mode, please follow the instructions given in Section 2.7, "Starting NFVD HA Platform".

To stop a component, or to get the status of a component on a given VM use the script":

- /opt/HPE/nfvd/bin/nfv-director.sh

```
•   # /opt/HPE/nfvd/bin/nfv-director.sh -h
•   Administration tool for the NFVD solution
•   Usage:
•     [options] [-c nfvdComponent] <action>
•     where action is one of start | stop | restart | status
•   options:
•      -c nfvdComponent : NFVD Component on which the action is applied
•   One of: activator | sosa | ecpool | lockmgr | openmediation | sitescope | uca-ebc | nfvd-agw |
    couchdb | uoc | idp | imageuploader
•   If not specified, the specified action applies to all installed NFVD components
•           -h               : Displays this usage message
•           -v               : Verbose mode
```

4.2 Maintenance Mode for NFVD HA VMs

A command line utility is provided that allows a particular NFVD VM to be put in maintenance mode.

Typical use cases to put a VM in maintenance mode would be

- Upgrade the VM with Operating System Patches
- Upgrade the VM with NFVD patches.
- Resizing the VM

When the particular VM is **put in maintenance**, the corresponding NFVD Service provided by that VM is stopped gracefully.

- In the case of Active/Active instances the other Active instances takes over the service provided by that VM
- In case of Active/Passive instance a switchover happens so the Passive VM becomes active to provide the service.

When the particular VM is put **out of maintenance** the services on that VM are started in case of Active/Active instances only.

The syntax of this command line utility would be

# /opt/HPE/nfvd/solutions/ha-example/maintenance_mode.sh start | stop | status

Puts the VM in maintenance mode.

where

start : Puts the specified node in maintenance

stop : Puts the specified node out of maintenance.

status : Indicates if the VM is in maintenance mode or not.

This command would be run by the user on the VM which needs to be put under maintenance.

The script will have the intelligence to identify which VM it is running on (GUI, FF, AA, SIS) and what service to stop/start exactly.

4.2.1 Maintenance of GUI VM

The command used on GUI VMs for maintenance operations is

```
[root@gui-node1]# /opt/HPE/nfvd/solutions/ha-example/maintenance_mode.sh start | stop | status
```

4.2.1.1 Example

gui-node1 maintenance mode status

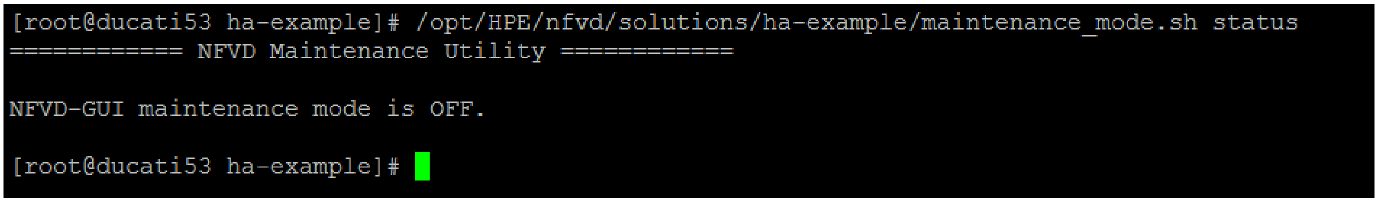

```
[root@ducati53 ha-example]# /opt/HPE/nfvd/solutions/ha-example/maintenance_mode.sh status
============ NFVD Maintenance Utility ============

NFVD-GUI maintenance mode is OFF.

[root@ducati53 ha-example]# 
```

Putting gui-node1 IN maintenance mode.

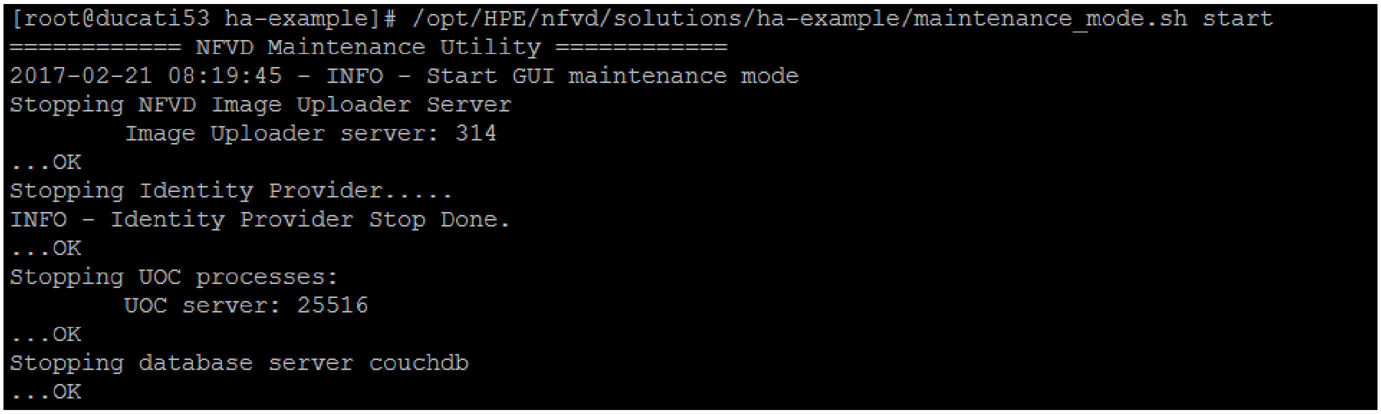

```
[root@ducati53 ha-example]# /opt/HPE/nfvd/solutions/ha-example/maintenance_mode.sh start
============ NFVD Maintenance Utility ============
2017-02-21 08:19:45 - INFO - Start GUI maintenance mode
Stopping NFVD Image Uploader Server
        Image Uploader server: 314
...OK
Stopping Identity Provider.....
INFO - Identity Provider Stop Done.
...OK
Stopping UOC processes:
        UOC server: 25516
...OK
Stopping database server couchdb
...OK
```

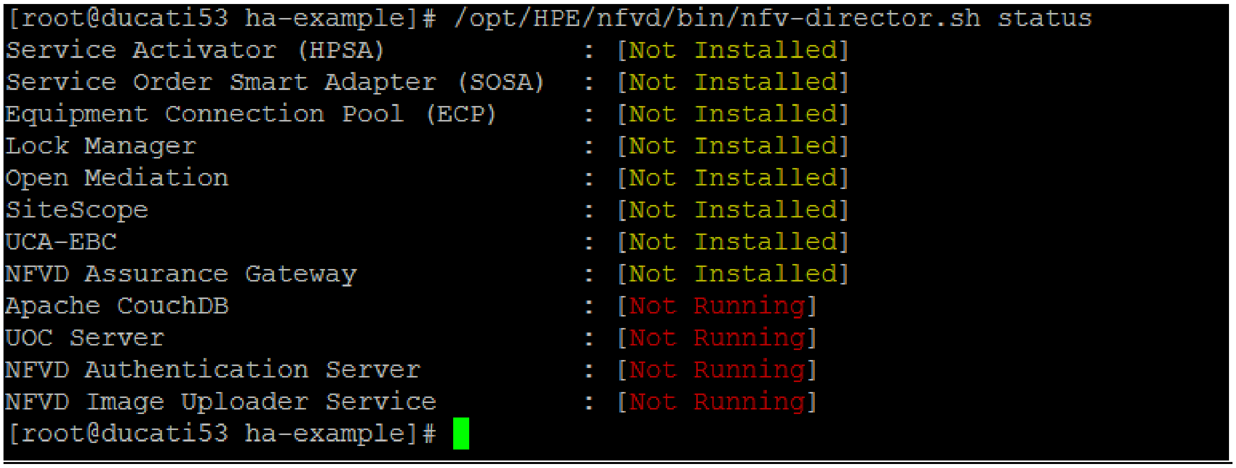

```
[root@ducati53 ha-example]# /opt/HPE/nfvd/bin/nfv-director.sh status
Service Activator (HPSA)              : [Not Installed]
Service Order Smart Adapter (SOSA)    : [Not Installed]
Equipment Connection Pool (ECP)       : [Not Installed]
Lock Manager                          : [Not Installed]
Open Mediation                        : [Not Installed]
SiteScope                             : [Not Installed]
UCA-EBC                               : [Not Installed]
NFVD Assurance Gateway                : [Not Installed]
Apache CouchDB                        : [Not Running]
UOC Server                            : [Not Running]
NFVD Authentication Server            : [Not Running]
NFVD Image Uploader Service           : [Not Running]
[root@ducati53 ha-example]#
```

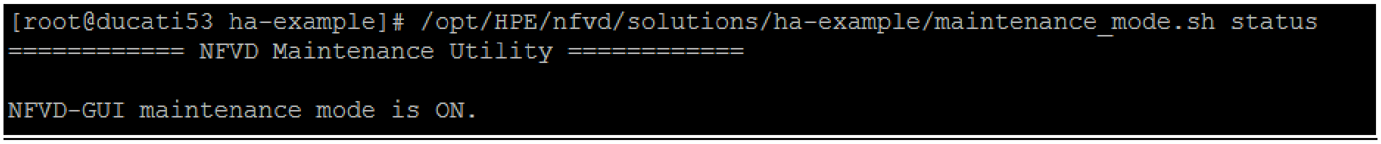

```
[root@ducati53 ha-example]# /opt/HPE/nfvd/solutions/ha-example/maintenance_mode.sh status
============ NFVD Maintenance Utility ============

NFVD-GUI maintenance mode is ON.
```

Taking gui-node1 OUT of maintenance mode.

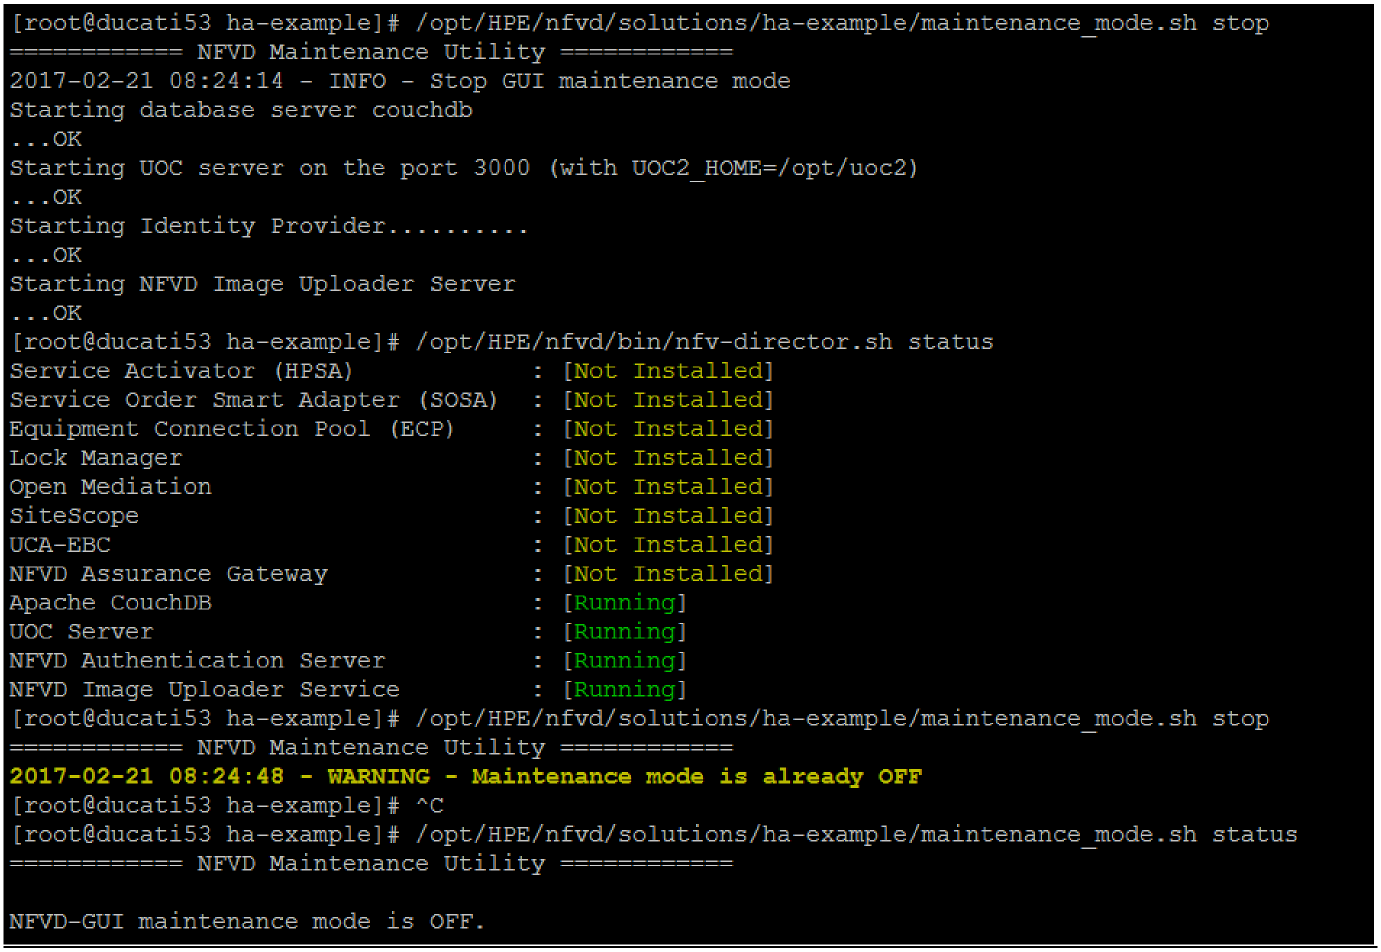

```
[root@ducati53 ha-example]# /opt/HPE/nfvd/solutions/ha-example/maintenance_mode.sh stop
============ NFVD Maintenance Utility ============
2017-02-21 08:24:14 - INFO - Stop GUI maintenance mode
Starting database server couchdb
...OK
Starting UOC server on the port 3000 (with UOC2_HOME=/opt/uoc2)
...OK
Starting Identity Provider..........
...OK
Starting NFVD Image Uploader Server
...OK
[root@ducati53 ha-example]# /opt/HPE/nfvd/bin/nfv-director.sh status
Service Activator (HPSA)              : [Not Installed]
Service Order Smart Adapter (SOSA)    : [Not Installed]
Equipment Connection Pool (ECP)       : [Not Installed]
Lock Manager                          : [Not Installed]
Open Mediation                        : [Not Installed]
SiteScope                             : [Not Installed]
UCA-EBC                               : [Not Installed]
NFVD Assurance Gateway                : [Not Installed]
Apache CouchDB                        : [Running]
UOC Server                            : [Running]
NFVD Authentication Server            : [Running]
NFVD Image Uploader Service           : [Running]
[root@ducati53 ha-example]# /opt/HPE/nfvd/solutions/ha-example/maintenance_mode.sh stop
============ NFVD Maintenance Utility ============
2017-02-21 08:24:48 - WARNING - Maintenance mode is already OFF
[root@ducati53 ha-example]# ^C
[root@ducati53 ha-example]# /opt/HPE/nfvd/solutions/ha-example/maintenance_mode.sh status
============ NFVD Maintenance Utility ============

NFVD-GUI maintenance mode is OFF.
```

4.2.2 Maintenance of Fulfilment VM

The command used on Fulfilment VMs for maintenance operations is

```
[root@ff-node1]# /opt/HPE/nfvd/solutions/ha-example/maintenance_mode.sh start | stop | status
```

4.2.2.1 Example

ff-node1 maintenance mode status

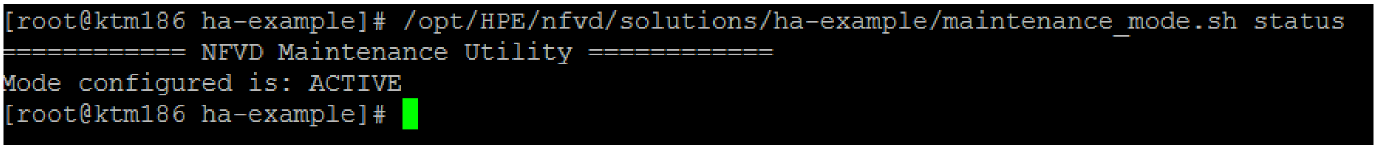

```
[root@ktm186 ha-example]# /opt/HPE/nfvd/solutions/ha-example/maintenance_mode.sh status
============ NFVD Maintenance Utility ============
Mode configured is: ACTIVE
[root@ktm186 ha-example]# 
```

Putting ff-node1 IN maintenance mode.

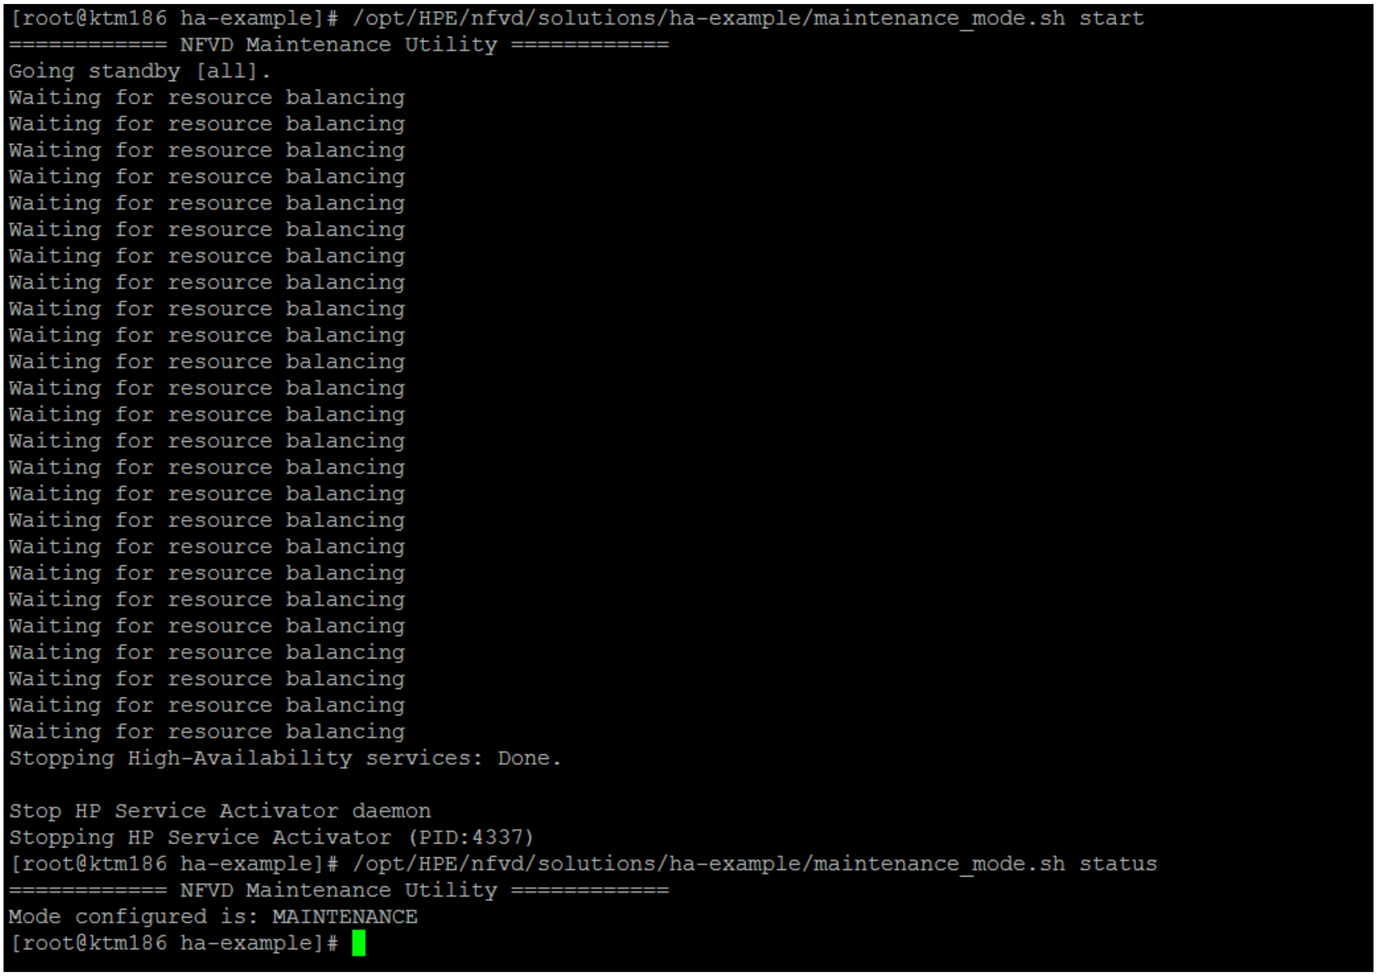

```
[root@ktm186 ha-example]# /opt/HPE/nfvd/solutions/ha-example/maintenance_mode.sh start
============ NFVD Maintenance Utility ============
Going standby [all].
Waiting for resource balancing
Waiting for resource balancing
Waiting for resource balancing
Waiting for resource balancing
Waiting for resource balancing
Waiting for resource balancing
Waiting for resource balancing
Waiting for resource balancing
Waiting for resource balancing
Waiting for resource balancing
Waiting for resource balancing
Waiting for resource balancing
Waiting for resource balancing
Waiting for resource balancing
Waiting for resource balancing
Waiting for resource balancing
Waiting for resource balancing
Waiting for resource balancing
Waiting for resource balancing
Waiting for resource balancing
Waiting for resource balancing
Waiting for resource balancing
Waiting for resource balancing
Waiting for resource balancing
Waiting for resource balancing
Stopping High-Availability services: Done.

Stop HP Service Activator daemon
Stopping HP Service Activator (PID:4337)
[root@ktm186 ha-example]# /opt/HPE/nfvd/solutions/ha-example/maintenance_mode.sh status
============ NFVD Maintenance Utility ============
Mode configured is: MAINTENANCE
[root@ktm186 ha-example]# 
```

Taking ff-node1 OUT of maintenance mode.

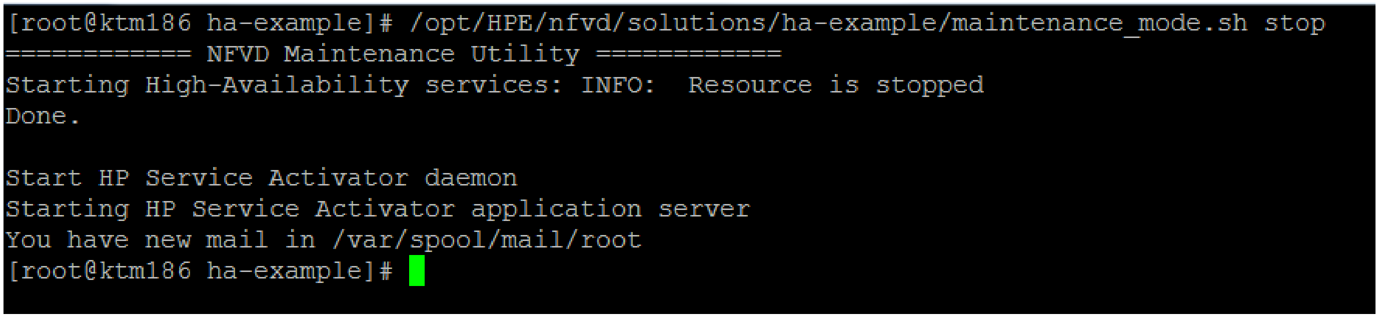

```
[root@ktm186 ha-example]# /opt/HPE/nfvd/solutions/ha-example/maintenance_mode.sh stop
============ NFVD Maintenance Utility ============
Starting High-Availability services: INFO:  Resource is stopped
Done.

Start HP Service Activator daemon
Starting HP Service Activator application server
You have new mail in /var/spool/mail/root
[root@ktm186 ha-example]# 
```

4.2.3 Maintenance Assurance VM

The command used on Assurance VMs for maintenance operations is

```
[root@aa-node1]# /opt/HPE/nfvd/solutions/ha-example/maintenance_mode.sh start | stop | status
```

4.2.3.1 Example

aa-node1 maintenance mode status

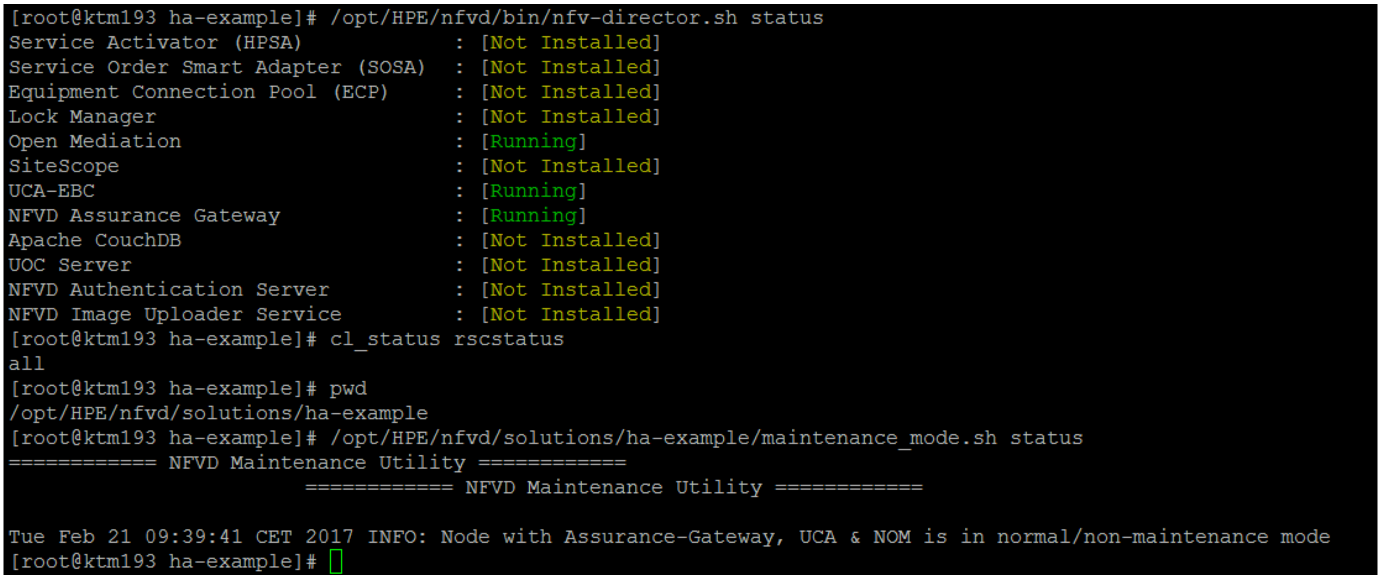

```
[root@ktm193 ha-example]# /opt/HPE/nfvd/bin/nfv-director.sh status
Service Activator (HPSA)            : [Not Installed]
Service Order Smart Adapter (SOSA)  : [Not Installed]
Equipment Connection Pool (ECP)     : [Not Installed]
Lock Manager                        : [Not Installed]
Open Mediation                      : [Running]
SiteScope                           : [Not Installed]
UCA-EBC                             : [Running]
NFVD Assurance Gateway              : [Running]
Apache CouchDB                      : [Not Installed]
UOC Server                          : [Not Installed]
NFVD Authentication Server          : [Not Installed]
NFVD Image Uploader Service         : [Not Installed]
[root@ktm193 ha-example]# cl_status rscstatus
all
[root@ktm193 ha-example]# pwd
/opt/HPE/nfvd/solutions/ha-example
[root@ktm193 ha-example]# /opt/HPE/nfvd/solutions/ha-example/maintenance_mode.sh status
============ NFVD Maintenance Utility ============
                        ============ NFVD Maintenance Utility ============

Tue Feb 21 09:39:41 CET 2017 INFO: Node with Assurance-Gateway, UCA & NOM is in normal/non-maintenance mode
[root@ktm193 ha-example]# 
```

Putting aa-node1 IN maintenance mode.

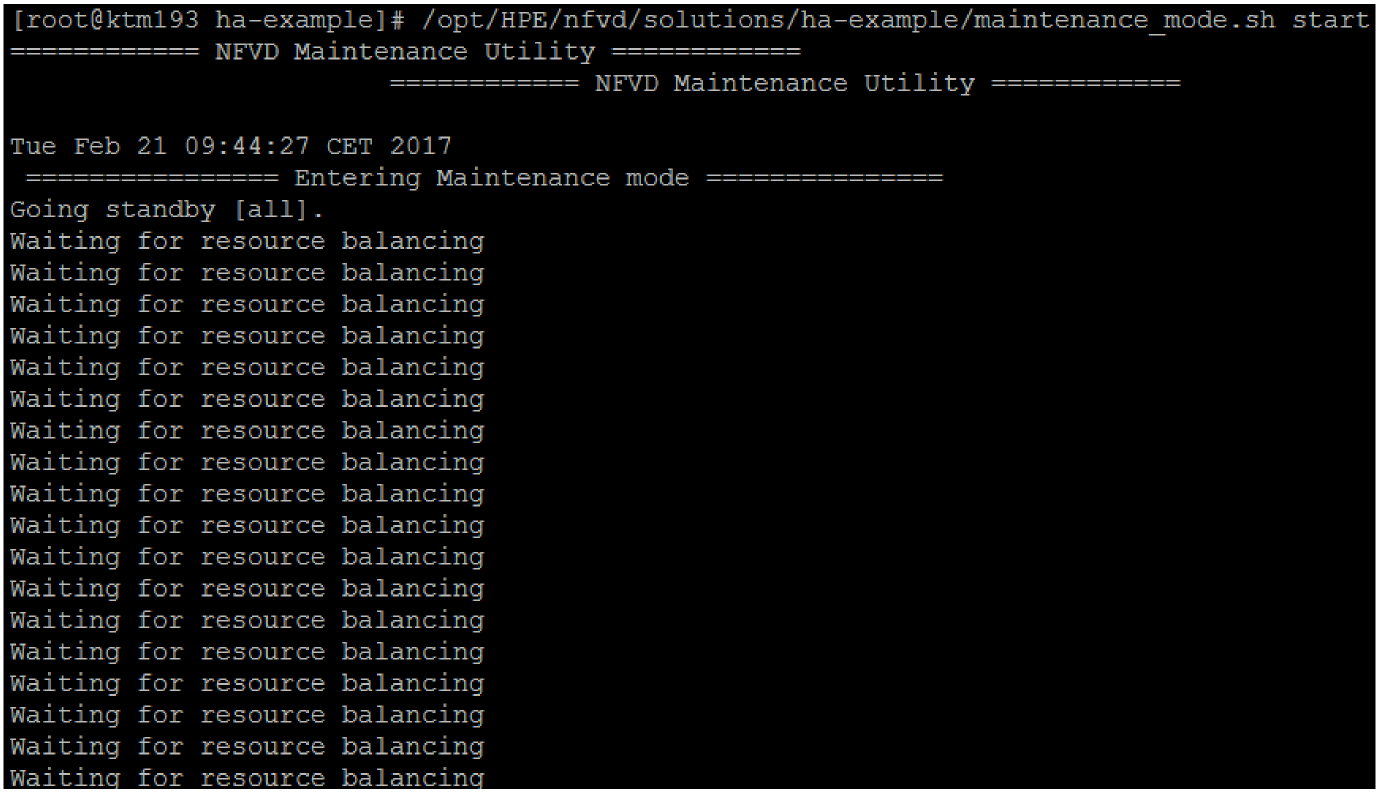

```
[root@ktm193 ha-example]# /opt/HPE/nfvd/solutions/ha-example/maintenance_mode.sh start
============ NFVD Maintenance Utility ============
                        ============ NFVD Maintenance Utility ============

Tue Feb 21 09:44:27 CET 2017
 ================ Entering Maintenance mode ===============
Going standby [all].
Waiting for resource balancing
Waiting for resource balancing
Waiting for resource balancing
Waiting for resource balancing
Waiting for resource balancing
Waiting for resource balancing
Waiting for resource balancing
Waiting for resource balancing
Waiting for resource balancing
Waiting for resource balancing
Waiting for resource balancing
Waiting for resource balancing
Waiting for resource balancing
Waiting for resource balancing
Waiting for resource balancing
Waiting for resource balancing
Waiting for resource balancing
Waiting for resource balancing
```

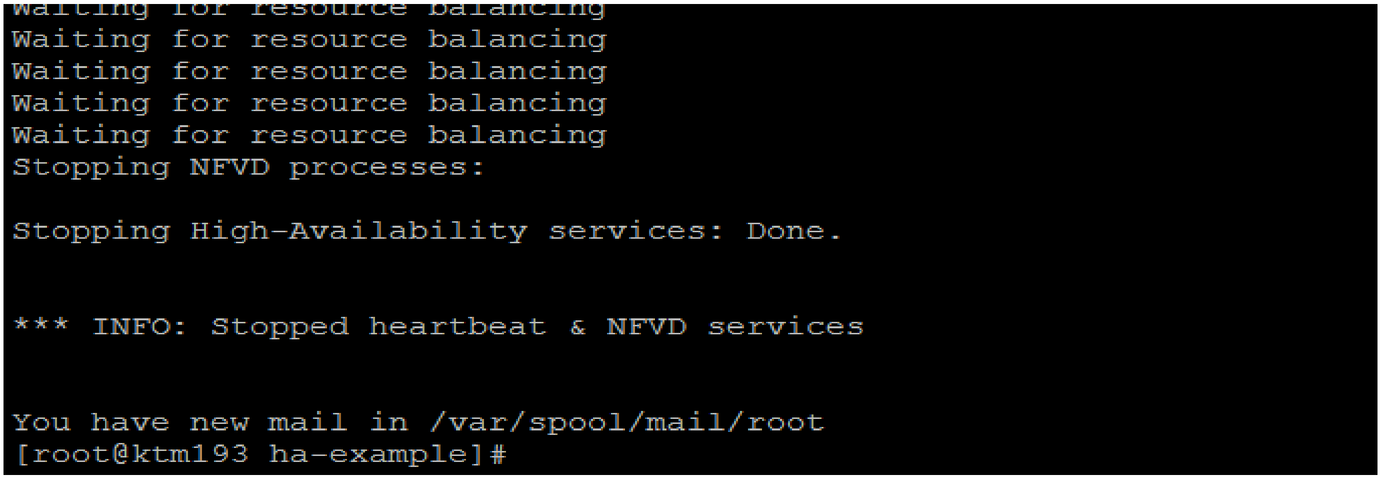

```
Waiting for resource balancing
Waiting for resource balancing
Waiting for resource balancing
Waiting for resource balancing
Waiting for resource balancing
Stopping NFVD processes:

Stopping High-Availability services: Done.


*** INFO: Stopped heartbeat & NFVD services


You have new mail in /var/spool/mail/root
[root@ktm193 ha-example]#
```

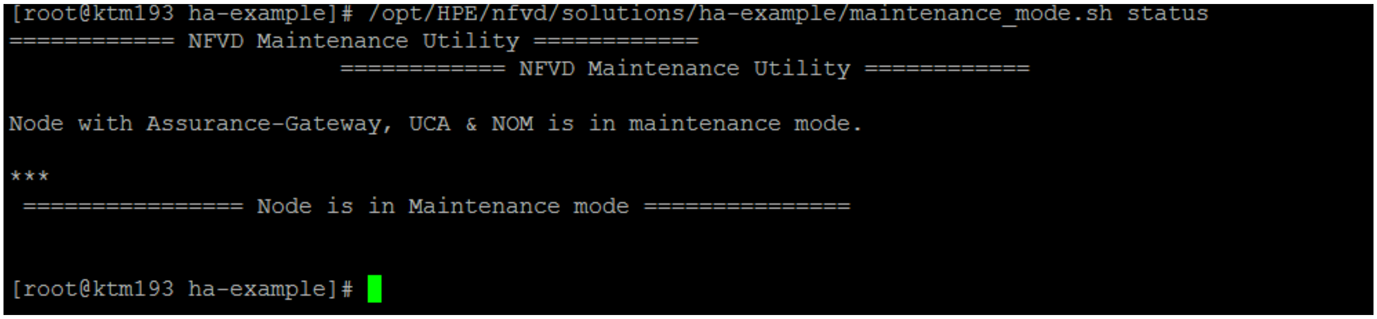

```
[root@ktm193 ha-example]# /opt/HPE/nfvd/solutions/ha-example/maintenance_mode.sh status
============ NFVD Maintenance Utility ============
                        ============ NFVD Maintenance Utility ============

Node with Assurance-Gateway, UCA & NOM is in maintenance mode.

***
 ================ Node is in Maintenance mode ===============


[root@ktm193 ha-example]#
```

Taking aa-node1 OUT of maintenance mode.

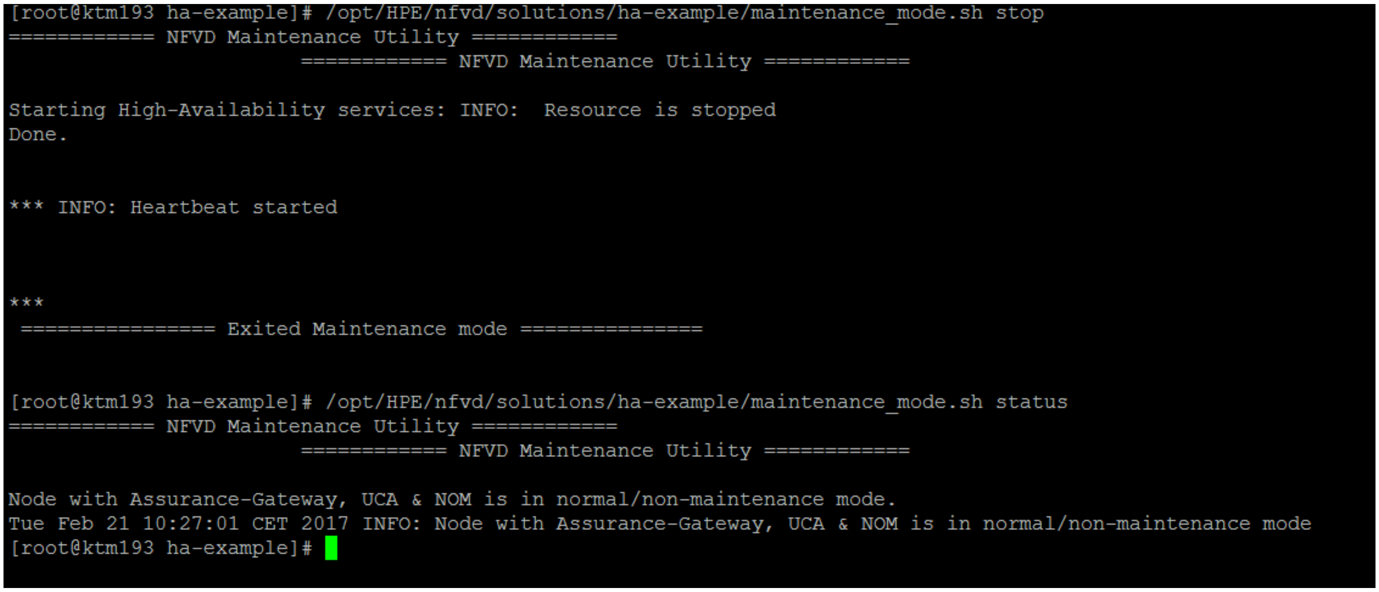

```
[root@ktm193 ha-example]# /opt/HPE/nfvd/solutions/ha-example/maintenance_mode.sh stop
============ NFVD Maintenance Utility ============
                        ============ NFVD Maintenance Utility ============

Starting High-Availability services: INFO:  Resource is stopped
Done.


*** INFO: Heartbeat started



***
 ================ Exited Maintenance mode ===============


[root@ktm193 ha-example]# /opt/HPE/nfvd/solutions/ha-example/maintenance_mode.sh status
============ NFVD Maintenance Utility ============
                        ============ NFVD Maintenance Utility ============

Node with Assurance-Gateway, UCA & NOM is in normal/non-maintenance mode.
Tue Feb 21 10:27:01 CET 2017 INFO: Node with Assurance-Gateway, UCA & NOM is in normal/non-maintenance mode
[root@ktm193 ha-example]# 
```

4.2.4 Maintenance SiteScope VM

The command used on SiteScope VMs for maintenance operations is

```
[root@sis-node1]# /opt/HPE/nfvd/solutions/ha-example/sis_maintenance_mode.sh start | stop | status
```

4.2.4.1 Example

sis-node1 maintenance mode status

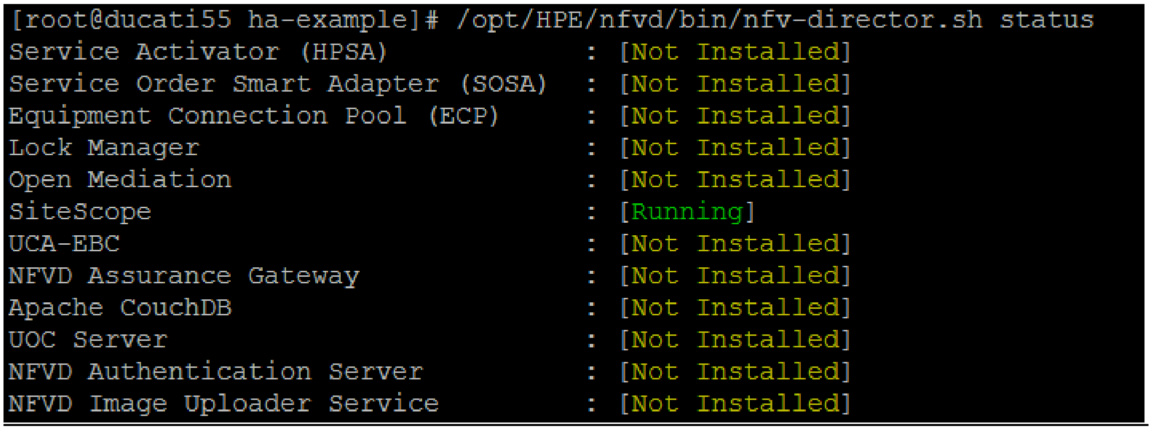

```
[root@ducati55 ha-example]# /opt/HPE/nfvd/bin/nfv-director.sh status
Service Activator (HPSA)              : [Not Installed]
Service Order Smart Adapter (SOSA)    : [Not Installed]
Equipment Connection Pool (ECP)       : [Not Installed]
Lock Manager                          : [Not Installed]
Open Mediation                        : [Not Installed]
SiteScope                             : [Running]
UCA-EBC                               : [Not Installed]
NFVD Assurance Gateway                : [Not Installed]
Apache CouchDB                        : [Not Installed]
UOC Server                            : [Not Installed]
NFVD Authentication Server            : [Not Installed]
NFVD Image Uploader Service           : [Not Installed]
```

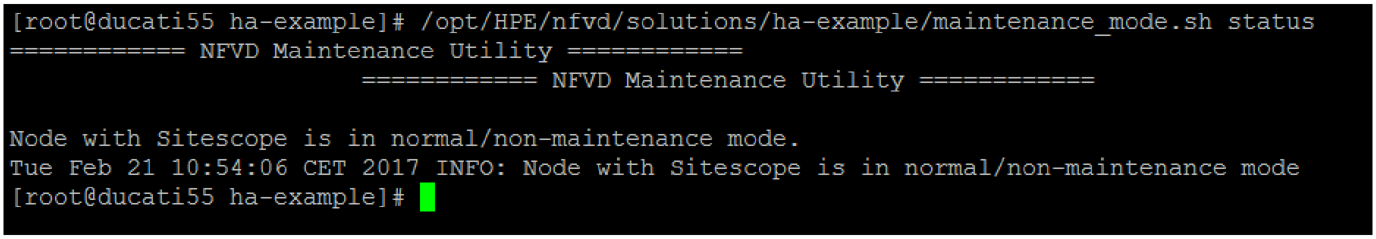

```
[root@ducati55 ha-example]# /opt/HPE/nfvd/solutions/ha-example/maintenance_mode.sh status
============ NFVD Maintenance Utility ============
                        ============ NFVD Maintenance Utility ============

Node with Sitescope is in normal/non-maintenance mode.
Tue Feb 21 10:54:06 CET 2017 INFO: Node with Sitescope is in normal/non-maintenance mode
[root@ducati55 ha-example]#
```

Putting sis-node1 IN maintenance mode.

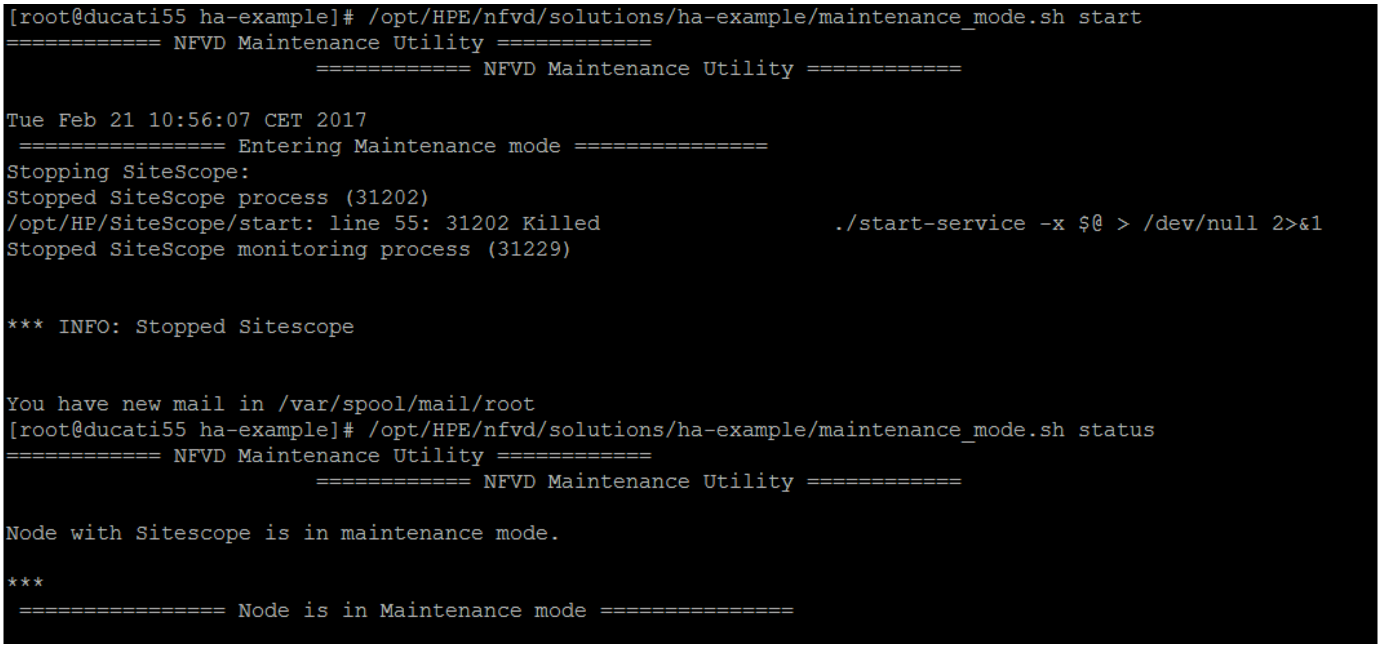

```
[root@ducati55 ha-example]# /opt/HPE/nfvd/solutions/ha-example/maintenance_mode.sh start
============ NFVD Maintenance Utility ============
                        ============ NFVD Maintenance Utility ============

Tue Feb 21 10:56:07 CET 2017
 ================ Entering Maintenance mode ===============
Stopping SiteScope:
Stopped SiteScope process (31202)
/opt/HP/SiteScope/start: line 55: 31202 Killed                  ./start-service -x $@ > /dev/null 2>&1
Stopped SiteScope monitoring process (31229)


*** INFO: Stopped Sitescope


You have new mail in /var/spool/mail/root
[root@ducati55 ha-example]# /opt/HPE/nfvd/solutions/ha-example/maintenance_mode.sh status
============ NFVD Maintenance Utility ============
                        ============ NFVD Maintenance Utility ============

Node with Sitescope is in maintenance mode.

***
 ================ Node is in Maintenance mode ===============
```

Taking sis-node1 OUT of maintenance mode.

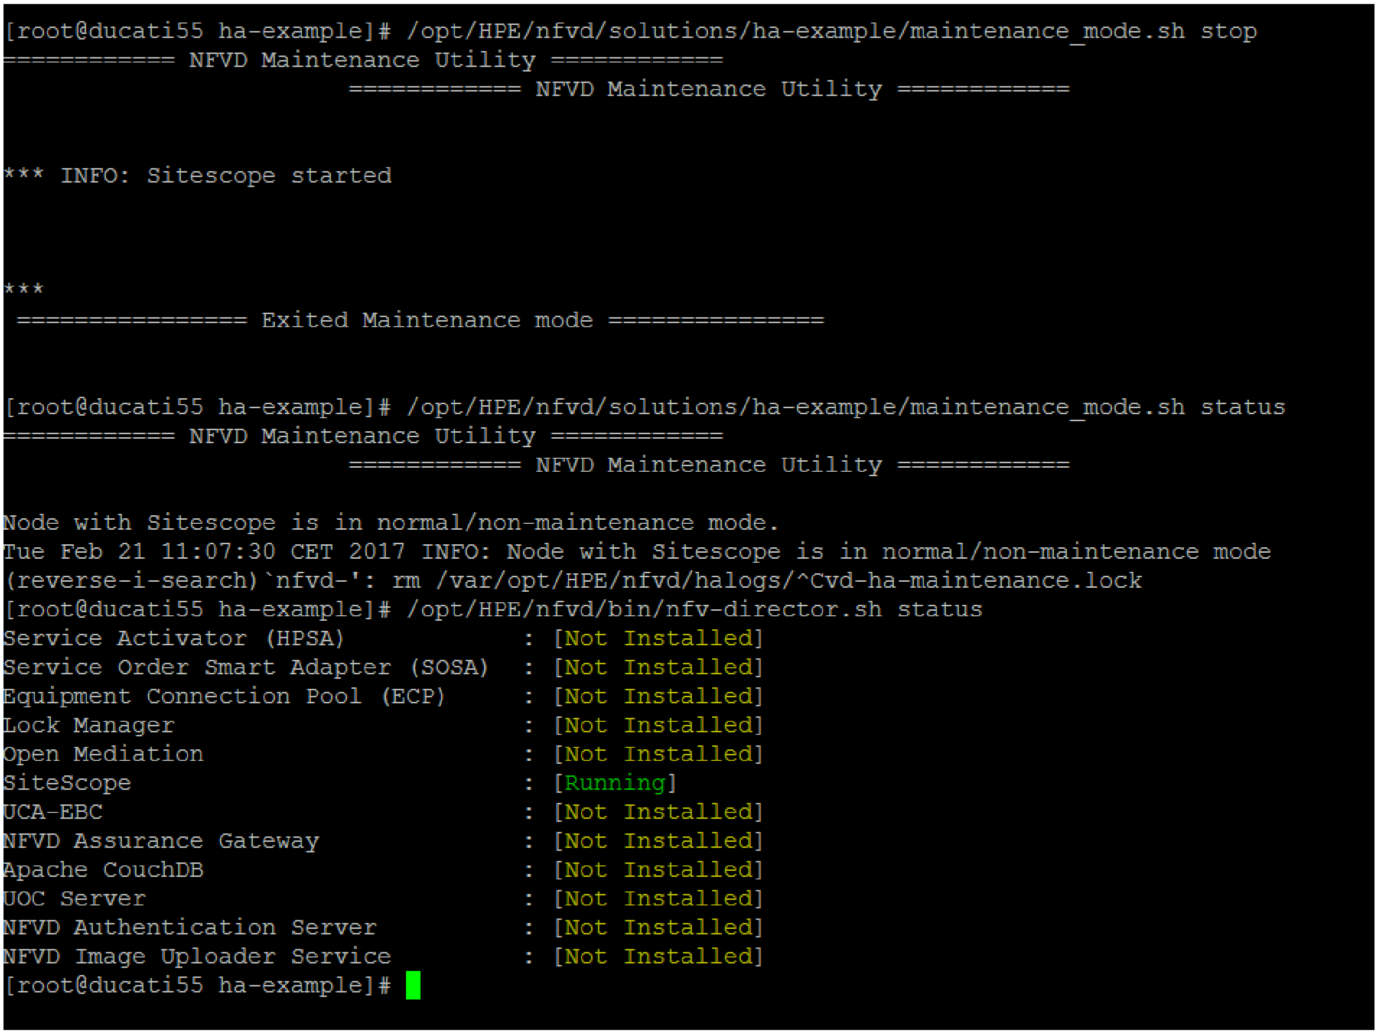

```
[root@ducati55 ha-example]# /opt/HPE/nfvd/solutions/ha-example/maintenance_mode.sh stop
============ NFVD Maintenance Utility ============
                        ============ NFVD Maintenance Utility ============


*** INFO: Sitescope started



***
 ================ Exited Maintenance mode ===============


[root@ducati55 ha-example]# /opt/HPE/nfvd/solutions/ha-example/maintenance_mode.sh status
============ NFVD Maintenance Utility ============
                        ============ NFVD Maintenance Utility ============

Node with Sitescope is in normal/non-maintenance mode.
Tue Feb 21 11:07:30 CET 2017 INFO: Node with Sitescope is in normal/non-maintenance mode
(reverse-i-search)`nfvd-': rm /var/opt/HPE/nfvd/halogs/^Cvd-ha-maintenance.lock
[root@ducati55 ha-example]# /opt/HPE/nfvd/bin/nfv-director.sh status
Service Activator (HPSA)            : [Not Installed]
Service Order Smart Adapter (SOSA)  : [Not Installed]
Equipment Connection Pool (ECP)     : [Not Installed]
Lock Manager                        : [Not Installed]
Open Mediation                      : [Not Installed]
SiteScope                           : [Running]
UCA-EBC                             : [Not Installed]
NFVD Assurance Gateway              : [Not Installed]
Apache CouchDB                      : [Not Installed]
UOC Server                          : [Not Installed]
NFVD Authentication Server          : [Not Installed]
NFVD Image Uploader Service         : [Not Installed]
[root@ducati55 ha-example]#
```

4.3 Backup / Restore

License files should be backed up as described in the “NFV Director Administration Guide”.

4.3.1 Backup / restore of Fulfilment data

Using Oracle tools you need to backup/restore all the database schema which has user as nfvd.

4.3.2 Backup / restore of GUI data

GUI does not persist and data hence no backup/restore of GUI data is not needed.

4.3.3 Backup / restore of Assurance data

On any of the assurance node you can use *uca-ebc-backup* tool to backup/restore Assurance data

As uca user

```
/opt/UCA-EBC/bin/uca-ebc-backup –h
This utility manages the backup/restore capabilities of UCA-EBC

Usage: uca-ebc-backup cmd [options], where cmd is one of following:

1. for a backup usage: uca-ebc-backup -backup [options]
   options are:
     -i instance     : specifies the UCA-EBC instance to backup [if not specified, default is
used]
     -from directory : specifies the data directory [default=/var/opt/UCA-EBC]
     -to directory   : specifies the backup directory [default=/var/opt/UCA-EBC/backup]
     -name filename  : specifies the filename for the backup [default=%instance-%date-%time]
     -excludelogs    : do not includes log files [if not specified, log files are backed up]

2. for a restore usage: uca-ebc-backup -restore -name filename [-to directory]
     -name filename  : specifies the backup filename (fully qualified) to restore
     -to directory   : specifies the data directory [default=/var/opt/UCA-EBC]

3. to list all available backup points: uca-ebc-backup -list [-from directory]
     -from directory : specifies the backup directory [default=/var/opt/UCA-EBC/backup]
```

4.3.4 Backup / Restore of SiteScope data

On SIS VM1 which is primary Sitescope node you can use tool *config_tool.sh* to backup/restore Sitescope data.

For backup select Export

```
# cd /opt/HP/SiteScope/bin/
# ./config_tool.sh -i console
This wizard enables you you to change the ports assigned to SiteScope,move configuration data
from one SiteScope installation to another.You can also configure an external agent for
integration with HP Operations Manager and BSM.

Select the actions that you want to perform.
--------------------------------------------------------------
--------------------------------------------------------------
Please select one of the options

->1 - Export: ()
  2 - Import: ()
  3 - Change ports: ()
  4 - HP Operations Agent: ()
```

For restore select Import

```
# /opt/HPE/nfvd/bin/nfv-director.sh –c sitescope stop
# /opt/HP/SiteScope/bin/config tool.sh -i console
./config_tool.sh -i console
This wizard enables you you to change the ports assigned to SiteScope,move configuration data
from one SiteScope installation to another.You can also configure an external agent for
integration with HP Operations Manager and BSM.

Select the actions that you want to perform.
--------------------------------------------------------------
--------------------------------------------------------------
Please select one of the options

->1 - Export: ()
  2 - Import: ()
  3 - Change ports: ()
  4 - HP Operations Agent: ()

# /opt/HPE/nfvd/bin/nfv-director.sh –c sitescope start
```

4.4 Upgrading an NFVD HA platform

If it is NOT possible to do a live update, then the NFVD system should be stopped during a maintenance window, and the installer script executed on both the primary and secondary nodes. The system can then be restarted.

It is important that the system being upgraded is able to run UCA-EBC. In an HA environment, UCA-EBC is an exclusive process that runs on one or the other platform.

Lets assume nfvdhaaa1 is Active (UCA running) and nfvdhaa2 is Passive and we want to upgrade nfvdhaaa1.

On nfvdhaaa1

```
touch /var/opt/HPE/nfvd/halogs/nfvd-ha-maintenance.lock
```

This will put nfvdhaaa1 in maintenance mode but it will not do a switchover to nfvdhaaa2.

We can then upgrade the platform which contains nfvdhaaa1.

Once the upgrade is done:

```
touch /var/opt/HPE/nfvd/halogs/nfvd-ha-maintenance.lock
```

Then perform the upgrade on the other platform by inverting the tasks. Always make sure when a platform is upgraded, that UCA-EBC is stopped on the other platform.

The order of the reconfiguration is important, you must start by upgrading/reconfiguring the secondary platform, then the primary platform. During this upgrade, it is essential to be sure that the platform being

Secondary platform:

```
# /opt/HPE/nfvd/install/nfvd-install.sh -s /kits/archives
```

Primary platform:

```
# /opt/HPE/nfvd/install/nfvd-install.sh -p /kits/archives
```

If you are upgrading an HA platform, you will need to tell the installer of that specificity using –p for the primary platform or –s for the secondary platform. If the currently installed version has no customized Definitions or Components, please execute the following command:

Secondary platform:

```
# /opt/HPE/nfvd/install/nfvd-install.sh -s -c <custom_files_dir> /kits/archives
```

Primary platform:

```
# /opt/HPE/nfvd/install/nfvd-install.sh -p -c <custom_files_dir> /kits/archives
```

Where

<custom_files_dir> is a directory containing customized files (Components, Definitions, Instances) replacing files from the kits. The directory content must match the same structure as in /opt/OV/ServiceActivator/solutions/NFVModel/etc/LoadXML

After these steps are performed, then using the reconfiguration tool will set the VIP to be used on both nodes, please use these exact commands:

Secondary platform:

```
# /opt/HPE/nfvd/install/nfvd-reconfigure.sh –s -f
```

Primary platform:

```
# /opt/HPE/nfvd/install/nfvd-reconfigure.sh -p -f
```

Chapter 5 Troubleshooting of NFVD HA VMs

Troubleshooting an HA installation of NFVD is very similar to troubleshooting a simplex installation; the reader should thus be familiar with the documents “NFV Director Administration Guide” and “NFV Director Troubleshooting Guide”.

There are a few extra considerations due to the clustering solutions used in the HA deployment; these are described in the following sections.

The first step in troubleshooting is to get the status of the various components on each node by running the nfv-director.sh status script:

```
opt/HPE/nfvd/bin/nfv-director.sh status
```

This will give a summary of the components running on the current node, as was described in

5.1 Problems with the GUI nodes

Normally, the two NFVD GUI nodes run in active/active mode, and are load-balanced by an external load balancer. The health of these nodes can be monitored by manually running the health check scripts:

```
curl -X GET  http://<GUI-node1-IP>:3001
curl -X GET  http://<GUI-node2-IP>:3001
```

These scripts should report success on BOTH GUI nodes. If only one of the nodes responds successfully, then NFVD is running in a degraded simplex mode and requires maintenance. The GUI troubleshooting procedures described in “NFV Director Troubleshooting Guide” should be followed for a failing node.

There also may be connectivity problems with the external load balancer. The above scripts should be run from the external load balancer or from a server on the same subnet as the load balancer in order to check connectivity.

5.2 Problems with the Fulfillment nodes

Fulfillment API processes

Normally, the two NFVD fulfillment API nodes run in active/active mode, and are load-balanced by an external load balancer. The health of these nodes can be monitored by manually running the health check scripts:

```
curl -X GET  http://<FF-node1-IP>:8080/nfvd-ext/status
curl -X GET  http://<FF-node2-IP>:8080/nfvd-ext/status
```

These scripts should report success on BOTH API nodes. If only one of the nodes responds successfully, then NFVD is running in a degraded simplex mode and requires maintenance. The API troubleshooting procedures described in “NFV Director Troubleshooting Guide” should be followed for a failing node.

There also may be connectivity problems with the external load balancer. The above scripts should be run from the external load balancer or from a server on the same subnet as the load balancer in order to check connectivity.

5.2.1 SOSA processes

A single instance of the SOSA processes runs on the Fulfillment nodes; if ever there is a problem with these processes or with the fulfillment node itself, then the Linux heartbeat daemon will move the processes to the other node.

However, if the SOSA processes are not running, then the system will not function correctly; the troubleshooting procedures described in “NFV Director Troubleshooting Guide” should be followed in this case.

To check if SOSA is running correctly, execute the following commands (assuming that ff-node1 is the active node):

```
 [root@ff-node1 ~]# cl_status rscstatus
```

The expected result is:

```
all
```

This confirms that ff-node1 is Active and is in control of resources

```
[root@ff-node1 ~]# /etc/init.d/activator-ep check
```

The expected result is:

```
HP Service Activator SOSA application is running
HP Service Activator Lock Manager application is running
```

5.2.2 Service Activator instances

There are two independent Service Activator instances running on the two fulfillment nodes. It can be verified that these two instances are running by logging in to each instance:

If only one instance is running, then NFVD is running in a degraded simplex mode and requires maintenance. The Fulfillment troubleshooting procedures described in “NFV Director Troubleshooting Guide” should be followed for a failing node.

Open a browser to the following address, this should present the Service Activator login page.

```
http://<ff-node1-IP>:8080/activator
```

5.3 Problems with Assurance nodes

A single instance of the Assurance processes runs on the Assurance nodes; if ever there is a problem with these processes or with the assurance node itself, then the Linux heartbeat daemon will move the processes to the other node.

However, if the Assurance processes are not running, then the system will not function correctly; the troubleshooting procedures described in “NFV Director Troubleshooting Guide” should be followed in this case.

To check if the Assurance node is running correctly, execute the following commands (assuming that aa-node1 is the active node):

On the Assurance VMs, execute the following command:

```
[root@aa-node1 ~]# cl_status rscstatus
```

The expected result is:

```
all
```

This confirms that aa-node1 is Active and is in control of resources

```
[root@aa-node1 ~]# /etc/init.d/nfvd status
```

The expected result is:

```
Status of NFVD processes
NFVD processes are running
```